
phyVIP-i.MX 8M Plus BSP Manual

DRAFT

PHYTEC Messtechnik GmbH

Sep 29, 2026

CONTENTS

1	1	PHYTEC Documentation	3
1.1	1.1	Supported Hardware	3
2	2	Getting Started	7
2.1	2.1	Get the Image	7
2.2	2.2	Write the Image to SD Card	7
2.3	2.3	First Start-up	10
3	3	Building the BSP	11
3.1	3.1	Basic Set-Up	11
3.2	3.2	Get the BSP	11
4	4	Installing the OS	15
4.1	4.1	Bootmode Switch (S1)	15
4.2	4.2	Flash eMMC	16
5	5	Development	23
5.1	5.1	Standalone Build preparation	23
5.2	5.2	U-Boot standalone build	24
5.3	5.3	Kernel standalone build	26
5.4	5.4	Host Network Preparation	27
5.5	5.5	Booting the Kernel from a Network	30
5.6	5.6	Accessing the Development states	31
5.7	5.7	Accessing the Latest Upstream Support	31
5.8	5.8	Format SD card	32
5.9	5.9	Ampliphy-boot	38
5.10	5.10	Working with FIT images	40
6	6	Device Tree (DT)	43
6.1	6.1	Introduction	43
6.2	6.2	PHYTEC i.MX 8M Plus BSP Device Tree Concept	43
7	7	Accessing Peripherals	49
7.1	7.1	i.MX 8M Plus Pin Muxing	49
7.2	7.2	RS232/RS485	50
7.3	7.3	Ethernet	53
7.4	7.4	WLAN/Bluetooth	56
7.5	7.5	SD card	57
7.6	7.6	eMMC Devices	57
7.7	7.7	SPI Master	64
7.8	7.8	GPIOs	65

7.9	7.9	I ² C Bus	67
7.10	7.10	EEPROM	67
7.11	7.11	RTC	69
7.12	7.12	USB Host Controller	71
7.13	7.13	CAN FD	71
7.14	7.14	Audio	73
7.15	7.15	Display	75
7.16	7.16	GPU	76
7.17	7.17	Power Management	76
7.18	7.18	Thermal Management	78
7.19	7.19	TPM	79
7.20	7.20	Watchdog	80
7.21	7.21	Power Key	81
7.22	7.22	AI Accelerators	82
7.23	7.23	ISP	83
7.24	7.24	On-Chip OTP Controller (OCOTP_CTRL) - eFuses	83

The table below shows the Compatible BSPs for this manual:

Compatible BSPs	BSP Type	Release	Yocto sion	Ver-	BSP Date	Release	BSP Status
BSP-Yocto-NXP-i.MX8MP-PD26.1.1	Minor		Walnascar		2026 q3		In Develop- ment

This BSP manual guides you through the installation and creation steps for the Board Support Package (BSP) and describes how to handle the interfaces for the **phyVIP-i.MX 8M Plus Vision Integration Platform**. Furthermore, this document describes how to create BSP images from the source code. This is useful for those who need to change the default image and need a way to implement these changes in a simple and reproducible way. Further, some sections of this manual require executing commands on a personal computer (host). Any and all of these commands are assumed to be executed on a Linux Operating System.

Note

This document contains code examples that describe the communication with the board over the serial shell. The code examples lines begin with `host:~$`, `target:~$` or `u-boot=>`. This describes where the commands are to be executed. Only after these keywords must the actual command be copied.

1 PHYTEC DOCUMENTATION

PHYTEC provides a variety of hardware and software documentation for all of its products. This includes any or all of the following:

QS Guide

A short guide on how to set up and boot a PHYTEC System-on-Module based board.

Hardware Manual

A detailed description of the System-on-Module and accompanying carrierboard.

Yocto Guide

A comprehensive guide for the Yocto version the phyCORE uses. This guide contains an overview of Yocto; introducing, installing, and customizing the PHYTEC BSP; how to work with programs like Poky and Bitbake; and much more.

BSP Manual

A manual specific to the BSP version of the phyCORE. Information such as how to build the BSP, booting, updating software, device tree, and accessing peripherals can be found here.

Development Environment Guide

This guide shows how to work with the Virtual Machine (VM) Host PHYTEC has developed and prepared to run various Development Environments. There are detailed step-by-step instructions for Eclipse and Qt Creator, which are included in the VM. There are instructions for running demo projects for these programs on a phyCORE product as well. Information on how to build a Linux host PC yourself is also a part of this guide.

Pin Muxing Table

phyCORE SOMs have an accompanying pin table (in Excel format). This table will show the complete default signal path, from the processor to the carrier board. The default device tree muxing option will also be included. This gives a developer all the information needed in one location to make muxing changes and design options when developing a specialized carrier board or adapting a PHYTEC phyCORE SOM to an application.

On top of these standard manuals and guides, PHYTEC will also provide Product Change Notifications, Application Notes, and Technical Notes. These will be done on a case-by-case basis. Most of the documentation can be found on the <https://www.phytec.eu/en/produkte/embedded-vision/phyvip-vision-integration-platform-imx-8m-plus/#downloads> of our product.

1.1 1.1 Supported Hardware

On our web page, you can see all supported Machines with the available Article Numbers for this release: BSP-Yocto-NXP-i.MX8MP-PD26.1.1 [download](#).

If you choose a specific **Machine Name** in the section **Supported Machines**, you can see which **Article Numbers** are available under this machine and also a short description of the hardware information. In case

you only have the **Article Number** of your hardware, you can leave the **Machine Name** drop-down menu empty and only choose your **Article Number**. Now it should show you the necessary **Machine Name** for your specific hardware

1.1.1 1.1.1 phyVIP-i.MX 8M Plus Components




Fig. 1: phyVIP (Hailo) Components (top)




Fig. 2: phyVIP (Hailo) Components (bottom)




Fig. 3: phyVIP (Kinara) Components (top)




Fig. 4: phyVIP (Kinara) Components (bottom)




Fig. 5: phyVIP PEB-B-006 expansion board Components (top)




Fig. 6: phyVIP PEB-B-006 expansion board Components (bottom)

2 GETTING STARTED

The **phyVIP-i.MX 8M Plus Vision Integration Platform** is shipped with a pre-flashed SD card. It contains the phytec-vision-image and can be used directly as a boot source. The eMMC is programmed with only a U-Boot by default. You can get all sources from the [BSP downloads](#) page. This chapter explains how to flash a BSP image to SD card and how to start the board.

There are several ways to flash an image to SD card or even eMMC. Most notably using simple, sequential writing with the Linux command line tool `dd`. An alternative way is to use PHYTEC's system initialization program called [partup](#), which makes it especially easy to format more complex systems. You can get [prebuilt Linux binaries of partup](#) from its release page. Also read [partup's README](#) for installation instructions.

2.1 2.1 Get the Image

The image contains all necessary files and makes sure partitions and any raw data are correctly written. Both the `partup` package and the WIC image, which can be flashed using `dd`, can be downloaded from our [BSP downloads](#) page.

Note that you can find different image versions and variants on our download server. The images are located on the server by folders per “BSP-Version”, “Distro-Name” and “Machine-Name”.

Example to download a `partup` package and a WIC image from the download server:

```
host:~$ wget https://download.phytec.de/Software/Linux/BSP-Yocto-i.MX8MP/BSP-Yocto-NXP-i.MX8MP-
↳ PD26.1.1/images/ampliphy-vendor/imx8mp-phyflex-phyvip-1/phytec-vision-image-imx8mp-phyflex-
↳ phyvip-1.rootfs.partup
host:~$ wget https://download.phytec.de/Software/Linux/BSP-Yocto-i.MX8MP/BSP-Yocto-NXP-i.MX8MP-
↳ PD26.1.1/images/ampliphy-vendor/imx8mp-phyflex-phyvip-1/phytec-vision-image-imx8mp-phyflex-
↳ phyvip-1.rootfs.wic.xz
```

Note

For eMMC, more complex partitioning schemes or even just large images, we recommend using the `partup` package, as it is faster in writing than `dd` and allows for a more flexible configuration of the target flash device.

2.2 2.2 Write the Image to SD Card

Warning

To create your bootable SD card, you must have root privileges on your Linux host PC. Be very careful when specifying the destination device! All files on the selected device will be erased immediately without any further query!

Selecting the wrong device may result in **data loss** and e.g. could erase your currently running system on your host PC!

2.2.1 2.2.1 Finding the Correct Device

To create your bootable SD card, you must first find the correct device name of your SD card and possible partitions. If any partitions of the SD cards are mounted, unmount those before you start copying the image to the SD card.

1. In order to get the correct device name, remove your SD card and execute:

```
host:~$ lsblk
```

2. Now insert your SD card and execute the command again:

```
host:~$ lsblk
```

3. Compare the two outputs to find the new device names listed in the second output. These are the device names of the SD card (device and partitions if the SD card was formatted).
4. In order to verify the device names being found, execute the command `sudo dmesg`. Within the last lines of its output, you should also find the device names, e.g. `/dev/sde` or `/dev/mmcblk0` (depending on your system).

Alternatively, you may use a graphical program of your choice, like [GNOME Disks](#) or [KDE Partition Manager](#), to find the correct device.

Now that you have the correct device name, e.g. `/dev/sde`, you can see the partitions which must be unmounted if the SD card is formatted. In this case, you will also find the device name with an appended number (e.g. `/dev/sde1`) in the output. These represent the partitions. Some Linux distributions automatically mount partitions when the device gets plugged in. Before writing, however, these need to be unmounted to avoid data corruption.

Unmount all those partitions, e.g.:

```
host:~$ sudo umount /dev/sde1
host:~$ sudo umount /dev/sde2
```

Now, the SD card is ready to be flashed with an image, using either `partup`, `dd` or `bmptool`.

2.2.2 2.2.2 Using bmptool

One way to prepare an SD card is using `bmptool`. Yocto automatically creates a block map file (`<IMAGENAME>-<MACHINE>.wic.bmap`) for the WIC image that describes the image content and includes checksums for data integrity. `bmptool` is packaged by various Linux distributions. For Debian-based systems install it by issuing:

```
host:~$ sudo apt install bmap-tools
```

Flash a WIC image to SD card by calling:

```
host:~$ bmptool copy phytec-vision-image-imx8mp-phyflex-phyvip-1?(.rootfs).wic?(.xz) /dev/<your_
↵device>
```

Replace <your_device> with your actual SD card's device name found previously, and make sure to place the file <IMAGENAME>-<MACHINE>.wic.bmap alongside the regular WIC image file, so bmaptool knows which blocks to write and which to skip.

Warning

bmaptool only overwrites the areas of an SD card where image data is located. This means that a previously written U-Boot environment may still be available after writing the image.

2.2.3 2.2.3 Using partup

Writing to an SD card with partup is done in a single command:

```
host:~$ sudo partup install phytec-vision-image-imx8mp-phyflex-phyvip-1?(.rootfs).partup /dev/  
↪<your_device>
```

Make sure to replace <your_device> with your actual device name found previously.

Further usage of partup is explained at its [official documentation website](#).

Warning

Host systems which are using resize2fs version 1.46.6 and older (e.g. Ubuntu 22.04) are not able to write partup packages created with Yocto Mickledore or newer to SD-Card. This is due to a new default option in resize2fs which causes an incompatibility. See [release notes](#).

Note

partup has the advantage of allowing to clear specific raw areas in the MMC user area, which is used in our provided partup packages to erase any existing U-Boot environments. This is a known issue *bmaptool* does not solve, as mentioned in the previous chapter.

Another key advantage of partup over other flashing tools is that it allows configuring MMC specific parts, like writing to eMMC boot partitions, without the need to call multiple other commands when writing.

2.2.4 2.2.4 Using dd

After having unmounted all SD card's partitions, you can create your bootable SD card.

Some PHYTEC BSPs produce uncompressed images (with filename-extension *.wic), and some others produce compressed images (with filename-extension *.wic.xz).

To flash an uncompressed images (*.wic) use command below:

```
host:~$ sudo dd if=phytec-vision-image-imx8mp-phyflex-phyvip-1?(.rootfs).wic of=/dev/<your_  
↪device> bs=1M conv=fsync status=progress
```

Or to flash a compressed images (*.wic.xz) use that command:

```
host:~$ xzcat phytec-vision-image-imx8mp-phyflex-phyvip-1?(.rootfs).wic.xz | sudo dd of=/dev/  
↪<your_device> bs=1M conv=fsync status=progress
```

Again, make sure to replace `<your_device>` with your actual device name found previously.

The parameter `conv=fsync` forces a sync operation on the device before `dd` returns. This ensures that all blocks are written to the SD card and none are left in memory. The parameter `status=progress` will print out information on how much data is and still has to be copied until it is finished.

2.3 2.3 First Start-up

- To boot from an SD card, the *bootmode switch (S1)* needs to be set to the following position:



- Insert the SD card
- Connect the target and the host with **USB to TTL serial adapter (1.8V)** on *(X8)* debug UART port
- Power up the board

3 BUILDING THE BSP

This section will guide you through the general build process of the i.MX 8M Plus BSP using Yocto and the phyLinux script. For more information about our meta-layer or Yocto in general visit: [Yocto Reference Manual \(walnascar\)](#).

3.1 Basic Set-Up

If you have never created a Phytec BSP with Yocto on your computer, you should take a closer look at the chapter BSP Workspace Installation in the Yocto Reference Manual ([walnascar](#)).

3.2 Get the BSP

There are two ways to get the BSP sources. You can download the complete BSP sources from our [BSP downloads](#) page; or you can fetch and build it yourself with Yocto. This is particularly useful if you want to make customizations.

The phyLinux script is a basic management tool for PHYTEC Yocto BSP releases written in Python. It is mainly a helper to get started with the BSP sources structure.

- Create a fresh project folder, get phyLinux, and make the script executable:

```
host:~$ mkdir ~/yocto
host:~$ cd yocto/
host:~/yocto$ wget https://download.phytec.de/Software/Linux/Yocto/Tools/phyLinux
host:~/yocto$ chmod +x phyLinux
```

Warning

A clean folder is important because phyLinux will clean its working directory. Calling phyLinux from a directory that isn't empty will result in a warning.

- Run phyLinux:

```
host:~/yocto$ ./phyLinux init
```

Note

On the first initialization, the phyLinux script will ask you to install the Repo tool in your `/usr/local/bin` directory.

- During the execution of the init command, you need to choose your processor platform (SoC), PHYTEC's BSP release number, and the hardware (MACHINE) you are working on.

Note

If you cannot identify your board with the information given in the selector, have a look at the invoice for the product. And have a look at the webpage of [our BSP](#).

- It is also possible to pass this information directly using command line parameters:

```
host:~/yocto$ DISTRO=ampliphy-vendor MACHINE=imx8mp-phyflex-phyvip-1 ./phyLinux init -p
↪ imx8mp -r BSP-Yocto-NXP-i.MX8MP-PD26.1.1
```

After the execution of the init command, phyLinux will print a few important notes. For example, it will print your git identity, SOC and BSP release which was selected as well as information for the next steps in the build process.

3.2.1 Starting the Build Process

- Set up the shell environment variables:

```
host:~/yocto$ source sources/poky/oe-init-build-env
```

Note

This needs to be done every time you open a new shell for starting builds.

- The current working directory of the shell should change to build/.
- Open the main configuration file and accept the GPU and VPU binary license agreements. Do this by uncommenting the corresponding line, as below.

```
host:~/yocto/build$ vim conf/local.conf
# Uncomment to accept NXP EULA
# EULA can be found under ../sources/meta-freescale/EULA
ACCEPT_FSL_EULA = "1"
```

- Build your image:

```
host:~/yocto/build$ bitbake phytec-vision-image
```

Note

For the first build we suggest starting with our smaller non-graphical image phytec-headless-image to see if everything is working correctly.

```
host:~/yocto/build$ bitbake phytec-headless-image
```

The first compile process takes about 40 minutes on a modern Intel Core i7. All subsequent builds will use the filled caches and should take about 3 minutes.

3.2.2 3.2.2 BSP Images

All images generated by Bitbake are deployed to `~/yocto/build/deploy*/images/<machine>`. The following list shows for example all files generated for the `imx8mp-phyflex-phyvip-1` machine:

- **u-boot.bin**: Binary compiled U-boot bootloader (U-Boot). Not the final Bootloader image!
- **oftree**: Default kernel device tree
- **u-boot-spl.bin**: Secondary program loader (SPL)
- **bl31-imx8mp.bin**: ARM Trusted Firmware binary
- **lpddr4_pmu_train_2d_dmem_202006.bin**, **lpddr4_pmu_train_2d_imem_202006.bin**: DDR PHY firmware images
- **imx-boot**: Bootloader build by `imx-mkimage` which includes SPL, U-Boot, ARM Trusted Firmware and DDR firmware. This is the final bootloader image which is bootable.
- **fitImage**: Linux kernel FIT image
- **fitImage-its*.its**
- **Image**: Linux kernel image
- **Image.config**: Kernel configuration
- **imx8mp-phyflex-phyvip*.dtb**: Kernel device tree file
- **imx8mp-phyflex-phyvip*.dtbo**: Kernel device tree overlay files
- **phytec-vision-image*.tar.gz**: Root file system
- **phytec-vision-image*.rootfs.wic.xz**: compressed SD card image
- ***.scr.uimg**: compiled bootscripts

4 INSTALLING THE OS

4.1 4.1 Bootmode Switch (S1)

Tip

Hardware revision baseboard: 1631.1 (Hailo), 1645.0 (Kinara)

The phyVIP features a boot switch with a single port to select the phyFLEX-i.MX 8M Plus FPSC Gamma default bootsource.

eMMC



Fig. 1: Boot from onboard e.MMC

SD card



Fig. 2: Boot from external SD card

Note

The phyVIP-i.MX 8M Plus additionally features a connector (*X12*) which can be utilized to select additional bootsources. For more information please refer to Hardware Manual and Schematics or contact

PHYTEC support.

To boot from e.MMC, make sure that the BSP image is flashed correctly to the e.MMC and the *bootmode switch* (*S1*) is set to **e.MMC**.

4.2 4.2 Flash e.MMC

For consistency, it is assumed that a TFTP server is configured; More importantly, all generated images, as listed above, are copied to the default /srv/tftp directory. If you do not have this set up, you need to adjust the paths that point to the images being used in the instructions. For instructions on how to set up the TFTP server and directory, see *Setup Network Host*.

4.2.1 4.2.1 Flash e.MMC from Network

i.MX 8M Plus boards have an Ethernet connector and can be updated over a network. Be sure to set up the development host correctly. The IP needs to be set to 192.168.3.10, the netmask to 255.255.255.0, and a TFTP server needs to be available. From a high-level point of view, an e.MMC device is like an SD card. Therefore, it is possible to flash the **WIC image** (<name>.wic) from the Yocto build system directly to the e.MMC. The image contains the bootloader, kernel, device tree, device tree overlays, and root file system.

4.2.1.1 Flash e.MMC via Network in Linux on Host

It is also possible to install the OS at e.MMC from your Linux host. As before, you need a complete image on your host.

Tip

A working network is necessary! *Setup Network Host*

Show your available image files on the host:

```
host:~$ ls /srv/tftp
phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.xz
phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.bmap
```

Send the image with the **bmaptool** command combined with **ssh** through the network to the e.MMC of your device:

```
host:~$ scp /srv/tftp/phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.* root@192.168.3.
↪11:/tmp && ssh root@192.168.3.11 "bmaptool copy /tmp/phytec-vision-image-imx8mp-phyflex-phyvip-
↪1.rootfs.wic.xz /dev/mmcblk2"
```

4.2.1.2 Flash e.MMC via Network in Linux on Target

You can update the e.MMC from your target.

Tip

A working network is necessary! *Setup Network Host*

Take a compressed or decompressed image with the accompanying block map file *.bmap on the host and send it with *ssh* through the network to the eMMC of the target with a one-line command:

```
target:~$ scp <USER>@192.168.3.10:/srv/tftp/phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.
↳wic.* /tmp && bmaptool copy /tmp/phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.xz /
↳dev/mmcblk2
```

4.2.1.3 Flash eMMC from Network in U-Boot on Target

These steps will show how to update the eMMC via a network.

Tip

This step only works if the size of the image file is less than 1GB due to limited usage of RAM size in the Bootloader after enabling OP-TEE.

Tip

A working network is necessary! *Setup Network Host*

Uncompress your image

```
host:~$ unxz /srv/tftp/phytec-headless-image-imx8mp-phyflex-phyvip-1.rootfs.wic.xz
```

Load your image via network to RAM:

- when using dhcp

```
u-boot=> dhcp phytec-headless-image-imx8mp-phyflex-phyvip-1.rootfs.wic
BOOTP broadcast 1
DHCP client bound to address 192.168.3.1 (1 ms)
Using ethernet@30be0000 device
TFTP from server 192.168.3.10; our IP address is 192.168.3.1
Filename 'phytec-headless-image-imx8mp-phyflex-phyvip-1.rootfs.wic'.
Load address: 0x40480000
Loading: #####
#####
#####
...
...
...
#####
#####
11.2 MiB/s
done
Bytes transferred = 911842304 (36599c00 hex)
```

- when using a static ip address (serverip and ipaddr must be set additionally).

```
u-boot=> tftp ${loadaddr} phytec-headless-image-imx8mp-phyflex-phyvip-1.rootfs.wic
Using ethernet@30be0000 device
TFTP from server 192.168.3.10; our IP address is 192.168.3.11
```

(continues on next page)

(continued from previous page)

```

Filename 'phytec-headless-image-imx8mp-phyflex-phyvip-1.rootfs.wic'.
Load address: 0x40480000
Loading: #####
          #####
          #####
          ...
          ...
          ...
          #####
          #####
          11.2 MiB/s
done
Bytes transferred = 911842304 (36599c00 hex)

```

Write the image to the eMMC:

```

u-boot=> mmc dev 2
switch to partitions #0, OK
mmc2(part 0) is current device
u-boot=> setexpr nblk ${filesize} / 0x200
u-boot=> mmc write ${loadaddr} 0x0 ${nblk}

MMC write: dev # 2, block # 0, count 1780942 ... 1780942 blocks written: OK

```

4.2.2 4.2.2 Flash e.MMC U-Boot image via Network from running U-Boot

Update the standalone U-Boot image `imx-boot` is also possible from U-Boot. This can be used if the bootloader on e.MMC is located in the e.MMC user area.

Tip

A working network is necessary! *Setup Network Host*

Load image over tftp into RAM and then write it to e.MMC:

```

u-boot=> tftp ${loadaddr} imx-boot
u-boot=> setexpr nblk ${filesize} / 0x200
u-boot=> mmc dev 2
u-boot=> mmc write ${loadaddr} 0x40 ${nblk}

```

Hint

The hexadecimal value represents the offset as a multiple of 512 byte blocks. See the *offset table* for the correct value of the corresponding SoC.

4.2.3 4.2.3 Flash e.MMC from USB stick

4.2.3.1 Flash e.MMC from USB in Linux

These steps will show how to flash the e.MMC on Linux with a USB stick. You only need a complete image saved on the USB stick and a bootable WIC image. (e.g. `phytec-vision-image-imx8mp-phyflex-phyvip-`

1. |yocto-imageext|). Set the *bootmode switch (S1)* to SD card.

- Insert and mount the USB stick:

```
[ 60.458908] usb-storage 1-1.1:1.0: USB Mass Storage device detected
[ 60.467286] scsi host0: usb-storage 1-1.1:1.0
[ 61.504607] scsi 0:0:0:0: Direct-Access                      8.07 PQ: 0 ANSI: 2
[ 61.515283] sd 0:0:0:0: [sda] 3782656 512-byte logical blocks: (1.94 GB/1.80 GiB)
[ 61.523285] sd 0:0:0:0: [sda] Write Protect is off
[ 61.528509] sd 0:0:0:0: [sda] No Caching mode page found
[ 61.533889] sd 0:0:0:0: [sda] Assuming drive cache: write through
[ 61.665969] sda: sda1
[ 61.672284] sd 0:0:0:0: [sda] Attached SCSI removable disk
target:~$ mount /dev/sda1 /mnt
```

- Now show your saved image files on the USB Stick:

```
target:~$ ls /mnt
phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.xz
phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.bmap
```

- Write the image to the phyFLEX-i.MX 8M Plus FPSC Gamma e.MMC (MMC device 2 without partition):

```
target:~$ bmaptool copy /mnt/phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.xz /dev/
↪mmcblk2
```

- After a complete write, your board can boot from e.MMC.

Tip

Before this will work, you need to configure the *bootmode switch (S1)* to eMMC.

4.2.3.2 Flash e.MMC from USB stick in U-Boot on Target

Tip

This step only works if the size of the image file is less than 1GB due to limited usage of RAM size in the Bootloader after enabling OPTEE.

These steps will show how to update the e.MMC via a USB device. Configure the *bootmode switch (S1)* to SD card and insert an SD card. Power on the board and stop in U-Boot prompt. Insert a USB device with the copied uncompressed WIC image to the USB slot.

Load your image from the USB device to RAM:

```
u-boot=> usb start
starting USB...
USB0:   USB EHCI 1.00
scanning bus 0 for devices... 2 USB Device(s) found
       scanning usb for storage devices... 1 Storage Device(s) found
u-boot=> fatload usb 0:1 0x58000000 phytec-headless-image-imx8mp-phyflex-phyvip-1.rootfs.wic
497444864 bytes read in 31577 ms (15 MiB/s)
```

Write the image to the eMMC:

```
u-boot=> mmc dev 2
switch to partitions #0, OK
mmc2(part 0) is current device
u-boot=> setexpr nblk ${filesize} / 0x200
u-boot=> mmc write 0x58000000 0x0 ${nblk}

MMC write: dev # 2, block # 0, count 1024000 ... 1024000 blocks written: OK
u-boot=> boot
```

4.2.4 4.2.4 Flash e.MMC from SD card

Even if there is no network available, you can update the eMMC. For that, you only need a ready-to-use image file (*.wic) located on the SD card. Because the image file is quite large, you need to allocate more SD card space. To create a new partition or enlarge your SD card, see [Resizing ext4 Root Filesystem](#).

Alternatively, flash a partup package to the SD card, as described in [Getting Started](#). This will ensure the full space of the SD card is used.

4.2.4.1 Flash e.MMC from SD card in Linux on Target

You can also flash the eMMC on Linux. You only need a partup package or WIC image saved on the SD card.

- Show your saved partup package or WIC image files on the SD card:

```
target:~$ ls
phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.partup
phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.xz
phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.bmap
```

- Write the image to the phyFLEX-i.MX 8M Plus FPSC Gamma e.MMC (MMC device 2 **without** partition) using [partup](#):

```
target:~$ partup install phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.partup /dev/
↪mmcblk2
```

Flashing the partup package has the advantage of using the full capacity of the eMMC device, adjusting partitions accordingly.

Note

Alternatively, `bmaptool` may be used instead:

```
target:~$ bmaptool copy phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.wic.xz /dev/
↪mmcblk2
```

Keep in mind that the root partition does not make use of the full space when flashing with `bmaptool`.

- After a complete write, your board can boot from eMMC.

Warning

Before this will work, you need to configure the *bootmode switch (S1)* to eMMC.

4.2.4.2 Flash e.MMC from SD card in U-Boot on Target

Tip

This step only works if the size of the image file is less than 1GB due to limited usage of RAM size in the Bootloader after enabling OPTEE. If the image file is too large use the *Updating e.MMC from SD card in Linux on Target* subsection.

- Flash an SD card with a working image and create a third ext4 partition. Copy the WIC image (for example phytec-vision-image.rootfs.wic) to this partition.
- Configure the *bootmode switch (S1)* to SD card and insert the SD card.
- Power on the board and stop in U-Boot.
- Load the image:

```
u-boot=> ext4load mmc 1:3 ${loadaddr} phytec-headless-image-imx8mp-phyflex-phyvip-1.rootfs.  
↪wic  
reading  
911842304 bytes read in 39253 ms (22.2 MiB/s)
```

- Switch the mmc dev to eMMC:

```
u-boot=> mmc list  
FSL_SDHC: 1 (SD)  
FSL_SDHC: 2 (eMMC)  
u-boot=> mmc dev 2  
switch to partitions #0, OK  
mmc2(part 0) is current device
```

- Flash your WIC image (for example phytec-vision-image.rootfs.wic) from the SD card to eMMC. This will partition the card and copy imx-boot, Image, dtb, dtbo, and root file system to eMMC.

```
u-boot=> setexpr nblk ${filesize} / 0x200  
u-boot=> mmc write ${loadaddr} 0x0 ${nblk}  
  
MMC write: dev # 2, block # 0, count 1780942 ... 1780942 blocks written: OK
```

- Power off the board and change the *bootmode switch (S1)* to eMMC.

5 DEVELOPMENT

5.1 5.1 Standalone Build preparation

In this section, we describe how to build the U-Boot and the Linux kernel without using the [Yocto Project](#). This procedure makes the most sense for development. The U-Boot source code, the Linux kernel, and all other git repositories are available on [GitHub](#).

5.1.1 5.1.1 Git Repositories

- Used U-Boot repository:

```
https://github.com/phytec/u-boot-phytec-imx
```

- Our U-Boot is based on the u-boot-phytec-imx and adds board-specific patches.
- Used Linux kernel repository:

```
https://github.com/phytec/linux-phytec-imx
```

- Our i.MX 8M Plus kernel is based on the linux-phytec-imx kernel.

To find out which u-boot and kernel tags to use for a specific board, have a look at your BSP source folder:

```
meta-phytec/recipes-kernel/linux/linux-phytec-imx_*.bb  
meta-phytec/recipes-bsp/u-boot/u-boot-phytec-imx_*.bb
```

5.1.2 5.1.2 Get the SDK

You can download the SDK from the [SDK downloads](#) page, or build it yourself with Yocto:

- Move to the Yocto build directory:

```
host:~$ source sources/poky/oe-init-build-env  
host:~$ bitbake -c populate_sdk phytec-vision-image # or another image
```

After a successful build the SDK installer is deployed to `build/deploy*/sdk`.

5.1.3 5.1.3 Install the SDK

- Set correct permissions and install the SDK:

```

host:~$ chmod +x phytec-ampliphy-vendor-glibc-x86_64-phytec-vision-image-cortexa53-crypto-
↳toolchain-5.2.4.sh
host:~$ ./phytec-ampliphy-vendor-glibc-x86_64-phytec-vision-image-cortexa53-crypto-
↳toolchain-5.2.4.sh
=====
Enter target directory for SDK (default: /opt/ampliphy-vendor/5.2.4):
You are about to install the SDK to "/opt/ampliphy-vendor/5.2.4". Proceed [Y/n]? Y
Extracting SDK...done
Setting it up...done
SDK has been successfully set up and is ready to be used.

```

5.1.4 5.1.4 Using the SDK

Activate the toolchain for your shell by sourcing the *environment-setup* file in the toolchain directory:

```
host:~$ source /opt/ampliphy-vendor/5.2.4/environment-setup-cortexa53-crypto-phytec-linux
```

5.1.5 5.1.5 Installing Required Tools

Building Linux and U-Boot out-of-tree requires some additional host tool dependencies to be installed. For Ubuntu you can install them with:

```
host:~$ sudo apt install bison flex libssl-dev
```

Warning

Using the SDK on older host distributions (e.g., Ubuntu 20.04 LTS) with Scarthgap NXP-based BSPs can cause issues when building U-Boot or Linux kernel tools for host use. If you encounter an “undefined reference” error, a workaround is to prepend the host’s binutils to the PATH.

```
host$ export PATH=/usr/bin:$PATH
```

Run this after sourcing the SDK *environment-setup* file.

Note, SDK issue has not been observed on newer distributions, such as Ubuntu 22.04, which appear to work without requiring any modifications.

5.2 5.2 U-Boot standalone build

5.2.1 5.2.1 Get the source code

- Get the U-Boot sources:

```
host:~$ git clone https://github.com/phytec/u-boot-phytec-imx
```

- To get the correct *U-Boot tag* you need to take a look at our release notes, which can be found here: [release notes](#)
- The **tag** used in this release is called v2025.04-2.2.0-phy
- Check out the needed *U-Boot tag*:

```
host:~$ cd ~/u-boot-phytec-imx/
host:~/u-boot-phytec-imx$ git fetch --all --tags
host:~/u-boot-phytec-imx$ git checkout tags/v2025.04-2.2.0-phy
```

- Set up a build environment:

```
host:~/u-boot-phytec-imx$ source /opt/ampliphy-vendor/5.2.4/environment-setup-cortexa53-
↳ crypto-phytec-linux
```

5.2.2 5.2.2 Get the needed binaries

To build the bootloader, you need to **copy** these **files** to your u-boot-phytec-imx **build directory** and rename them to fit with *mkimage* script:

- **ARM Trusted firmware binary** (*mkimage tool* compatible format **bl31.bin**): bl31-imx8mp.bin
- **OPTEE image** (optional): tee.bin
- **DDR firmware files** (*mkimage tool* compatible format **lpddr4_[i,d]mem_*d_*.bin**):
 lpddr4_dmem_1d_*.bin, lpddr4_dmem_2d_*.bin, lpddr4_imem_1d_*.bin,
 lpddr4_imem_2d_*.bin

If you already built our BSP with Yocto, you can get the bl31-imx8mp.bin, tee.bin and lpddr4_*.bin from the directory mentioned here: *BSP Images*

Or you can download the files here: <https://download.phytec.de/Software/Linux/BSP-Yocto-i.MX8MP/BSP-Yocto-NXP-i.MX8MP-PD26.1.1/images/ampliphy-vendor/imx8mp-phyflex-phyvip-1/imx-boot-tools/>

Warning

Make sure you rename the files you need so that they are compatible with the *mkimage tool*.

5.2.3 5.2.3 Build the bootloader

- build flash.bin (imx-boot):

```
host:~/u-boot-phytec-imx$ make imx8mp-phyflex-phyvip_defconfig
host:~/u-boot-phytec-imx$ make flash.bin
```

5.2.4 5.2.4 Flash the bootloader to a block device

The flash.bin can be found at u-boot-phytec-imx/ directory and now can be flashed. A chip-specific offset is needed:

SoC	Offset User Area	Offset Boot Partition	e.MMC Device
i.MX 8M Plus	32 kiB	0 kiB	/dev/mmcblk2

E.g. flash SD card:

```
host:~/u-boot-phytec-imx$ sudo dd if=flash.bin of=/dev/sd[x] bs=1024 seek=32 conv=fsync
```

Hint

The specific offset values are also declared in the Yocto variables “BOOTLOADER_SEEK” and “BOOTLOADER_SEEK_EMMC”

5.3 5.3 Kernel standalone build

The kernel is packaged in a FIT image together with the device tree. U-Boot has been adapted to be able to load a FIT image and boot the kernel contained in it. As a result, the kernel Image has to be packaged in a FIT image.

5.3.1 5.3.1 Setup sources

- The used linux-phytec-imx branch can be found in the [release notes](#)
- The tag needed for this release is called v6.12.49-2.2.0-phy
- Check out the needed linux-phytec-imx tag:

```
host:~$ git clone https://github.com/phytec/linux-phytec-imx
host:~$ cd ~/linux-phytec-imx/
host:~/linux-phytec-imx$ git fetch --all --tags
host:~/linux-phytec-imx$ git checkout tags/v6.12.49-2.2.0-phy
```

- For committing changes, it is highly recommended to switch to a new branch:

```
host:~/linux-phytec-imx$ git switch --create <new-branch>
```

- Set up a build environment:

```
host:~/linux-phytec-imx$ source /opt/ampliphy-vendor/5.2.4/environment-setup-cortexa53-
➔ crypto-phytec-linux
```

5.3.2 5.3.2 Build the kernel

- Build the linux kernel:

```
host:~/linux-phytec-imx$ make imx8_phytec_defconfig
host:~/linux-phytec-imx$ make -j$(nproc)
```

- Install kernel modules to e.g. NFS directory:

```
host:~/linux-phytec-imx$ make INSTALL_MOD_PATH=/home/<user>/<rootfspath> modules_install
```

- The Image can be found at ~/linux-phytec-imx/arch/arm64/boot/Image.gz
- The dtb can be found at ~/linux-phytec-imx/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip.dtb
- For (re-)building only Devicetrees and -overlays, it is sufficient to run

```
host:~/linux-phytec-imx$ make dtbs
```

or, to build a specific dtb (e.g. imx8mp-phyflex-phyvip.dtb):

```
host:~/linux-phytec-imx$ make freescale/imx8mp-phyflex-phyvip.dtb
```

Note

If you are facing the following build issue:

```
scripts/dtc/yamltree.c:9:10: fatal error: yaml.h: No such file or directory
```

Make sure you installed the package “*libyaml-dev*” on your host system:

```
host:~$ sudo apt install libyaml-dev
```

5.3.3 5.3.3 Package the kernel in a FIT image

To simply replace the kernel, you will need an `image tree source (.its)` file. If you already built our BSP with Yocto, you can get the its file from the directory mentioned here: *BSP Images* Or you can download the file here: <https://download.phytec.de/Software/Linux/BSP-Yocto-i.MX8MP/BSP-Yocto-NXP-i.MX8MP-PD26.1.1/images/ampliphy-vendor/imx8mp-phyflex-phyvip-1/>

Copy the .its file to the current working directory, create a link to the kernel image and create the final fitImage with mkimage.

```
host:~/linux-phytec-imx$ cp /path/to/yocto/deploydir/fitimage-its*.its .
&& ln -s arch/arm64/boot/Image.gz linux.bin
&& mkimage -f fitImage-its*.its fitImage
```

5.3.4 5.3.4 Copy FIT image and kernel modules to SD card

When one-time boot via netboot is not sufficient, the FIT image along with the kernel modules may be copied directly to a mounted SD card.

```
host:~/linux-phytec-imx$ cp fitImage /path/to/sdcard/boot/
host:~/linux-phytec-imx$ make INSTALL_MOD_PATH=/path/to/sdcard/root/ modules_install
```

5.4 5.4 Host Network Preparation

For various tasks involving a network in the Bootloader, some host services are required to be set up. On the development host, a TFTP, NFS and DHCP server must be installed and configured. The following tools will be needed to boot via Ethernet:

```
host:~$ sudo apt install tftpd-hpa nfs-kernel-server kea
```

5.4.1 5.4.1 TFTP Server Setup

- First, create a directory to store the TFTP files:

```
host:~$ sudo mkdir /srv/tftp
```

- Then copy your BSP image files to this directory and make sure other users have read access to all the files in the tftp directory, otherwise they are not accessible from the target.

```
host:~$ sudo chmod -R o+r /srv/tftp
```

- You also need to configure a static IP address for the appropriate interface. The default IP address of the PHYTEC evaluation boards is 192.168.3.11. Setting a host address 192.168.3.10 with netmask 255.255.255.0 is a good choice.

```
host:~$ ip addr show <network-interface>
```

Replace <network-interface> with the network interface you configured and want to connect the board to. You can show all network interfaces by not specifying a network interface.

- The message you receive should contain this:

```
inet 192.168.3.10/24 brd 192.168.3.255
```

- Create or edit the /etc/default/tftpd-hpa file:

```
# /etc/default/tftpd-hpa

TFTP_USERNAME="tftp"
TFTP_DIRECTORY="/srv/tftp"
TFTP_ADDRESS=":69"
TFTP_OPTIONS="-s -c"
```

- Set TFTP_DIRECTORY to your TFTP server root directory
- Set TFTP_ADDRESS to the host address the server is listening to (set to 0.0.0.0:69 to listen to all local IPs)
- Set TFTP_OPTIONS, the following command shows the available options:

```
host:~$ man tftpd
```

- Restart the services to pick up the configuration changes:

```
host:~$ sudo service tftpd-hpa restart
```

Now connect the ethernet port of the board to your host system. We also need a network connection between the embedded board and the TFTP server. The server should be set to IP 192.168.3.10 and netmask 255.255.255.0.

5.4.1.1 NFS Server Setup

- Create an nfs directory:

```
host:~$ sudo mkdir /srv/nfs
```

- Temporarily export the nfs directory: The NFS server is not restricted to a certain file system location, so all we have to do is to export our root file system to the embedded network. In this example, the whole directory is exported and the “lab network” address of the development host is 192.168.3.10. The IP address has to be adapted to the local needs:

```
host:~$ sudo exportfs -i -o rw,no_root_squash,sync,no_subtree_check 192.168.3.0/
↪255.255.255.0:/srv/nfs
```

- unexport the rootfs when finished:

```
host:~$ sudo exportfs -u 192.168.3.0/255.255.255.0:/srv/nfs
```

5.4.1.1.1 Permanent export

- To make the export persistent across reboots on most distributions, modify the `/etc/exports` file and export it:

```
/srv/nfs 192.168.3.0/255.255.255.0(rw,no_root_squash,sync,no_subtree_check)
```

- Now the NFS-Server has to read the `/etc/exportfs` file again:

```
host:~$ sudo exportfs -ra
```

5.4.1.2 DHCP Server setup

- Create or edit the `/etc/kea/kea-dhcp4.conf` file; Using the internal subnet sample. Replace `<network-interface>` with the name for the physical network interface:

```
{
  "Dhcp4": {
    "interfaces-config": {
      "interfaces": [ "<network-interface>/192.168.3.10" ]
    },
    "lease-database": {
      "type": "memfile",
      "persist": true,
      "name": "/tmp/dhcp4.leases"
    },
    "valid-lifetime": 28800,
    "subnet4": [{
      "id": 1,
      "next-server": "192.168.3.10",
      "subnet": "192.168.3.0/24",
      "pools": [
        { "pool": "192.168.3.1 - 192.168.3.255" }
      ]
    }]
  }
}
```

Warning

Be careful when creating subnets as this may interfere with the company network policy. To be on the safe side, use a different network and specify that via the `interfaces` configuration option.

- Now the DHCP-Server has to read the `/etc/kea/kea-dhcp4.conf` file again:

```
host:~$ sudo systemctl restart kea-dhcp4-server
```

When you boot/restart your host PC and don't have the network interface, as specified in the `kea-dhcp4` config, already active the `kea-dhcp4-server` will fail to start. Make sure to start/restart the `systemd` service when you connect the interface.

Note

DHCP server setup is only needed when using dynamic IP addresses. For our vendor BSPs, static IP addresses are used by default.

```
u-boot=> env print ip_dyn
ip_dyn=no
```

To use dynamic IP addresses for netboot, `ip_dyn` needs to be set to `yes`.

5.5 5.5 Booting the Kernel from a Network

Bootting from a network means loading the kernel and device tree over TFTP and the root file system over NFS. The bootloader itself must already be loaded from another available boot device.

5.5.1 5.5.1 Place Images on Host for Netboot

- Copy the kernel fitimage to your tftp directory:

```
host:~$ cp fitImage /srv/tftp
```

- Copy the bootscript to your tftp directory:

```
host:~$ cp net_boot_fit.scr.uimg /srv/tftp/
```

- Make sure other users have read access to all the files in the tftp directory, otherwise they are not accessible from the target:

```
host:~$ sudo chmod -R o+r /srv/tftp
```

- Extract the rootfs to your nfs directory:

```
host:~$ sudo tar -xvzf phytec-vision-image-imx8mp-phyflex-phyvip-1.rootfs.tar.gz -C /srv/nfs
```

Note

Make sure you extract with `sudo` to preserve the correct ownership.

5.5.2 5.5.2 Network Settings on Target

To customize the targets ethernet configuration, please follow the description here: *Network Environment Customization*

5.5.3 5.5.3 Booting from an Embedded Board

- Boot the board into the U-boot prompt and press any key to hold.
- To boot from a network, call:

```
u-boot=> setenv boot_targets ethernet
u-boot=> bootflow scan -lb
```

- To persistently boot from network, save the environment after setting `boot_targets` to `ethernet`.

```
u-boot=> env set boot_targets "ethernet"
u-boot=> env save
```

- To use DHCP for booting from a network:

```
u-boot=> env set ip_dyn true
u-boot=> env save
```

- When not using DHCP (`ip_dyn` set to false), which is the default for our NXP BSPs, U-Boot needs `ipaddr`, `serverip` and `nfs_eth_dev` to be set. `nfs_eth_dev` describes the name of the ethernet interface in Linux before any renaming to the endx naming scheme is applied. By default this should be set to the SoM Ethernet in the board environment. If values different from the PHYTEC provided defaults are desired, these variables can be set in the U-Boot environment:

```
u-boot=> env set ipaddr <xxx.xxx.xxx.xxx>
u-boot=> env set serverip <xxx.xxx.xxx.xxx>
u-boot=> env set nfs_eth_dev eth<x>
u-boot=> env save
```

5.6 5.6 Accessing the Development states

5.6.1 5.6.1 Development state of current release

These release manifests exist to give you access to the development states of the *Yocto* BSP. They will not be displayed in the phyLinux selection menu but need to be selected manually. This can be done using the following command line:

```
host:~$ ./phyLinux init -p imx8mp -r BSP-Yocto-NXP-i.MX8MP-PD26.1.y
```

This will initialize a BSP that will track the latest development state of the current release (BSP-Yocto-NXP-i.MX8MP-PD26.1.1). From now on *repo sync* in this folder will pull all the latest changes from our Git repositories:

```
host:~$ repo sync
```

5.6.2 5.6.2 Development state of upcoming release

Also development states of upcoming releases can be accessed this way. For this execute the following command and look for a release with a higher PDXX.Y number than the latest one (BSP-Yocto-NXP-i.MX8MP-PD26.1.1) and `.y` at the end:

```
host:~$ ./phyLinux init -p imx8mp
```

5.7 5.7 Accessing the Latest Upstream Support

We have a vanilla manifest that makes use of the Yocto master branches (not an NXP release), Linux, and U-Boot. This can be used to test the latest upstream kernel/U-Boot.

Note

The master manifest reflects the latest state of development. This tends to be broken from time to time. We try to fix the master on a regular basis.

```
host:~$ ./phyLinux init -p imx8mp -r BSP-Yocto-Ampliphy-i.MX8MP-master
```

5.8 5.8 Format SD card

Most images are larger than the default root partition. To flash any storage device with SD Card, the rootfs needs to be expanded or a separate partition needs to be created. There are some different ways to format the SD Card. The easiest way to do this is to use the UI program Gparted.

5.8.1 5.8.1 Gparted

- Get GParted:


```
host:~$ sudo apt install gparted
```

- Insert the SD card into your host and get the device name:

```
host:~$ dmesg | tail
...
[30436.175412] sd 4:0:0:0: [sdb] 62453760 512-byte logical blocks: (32.0 GB/29.8 GiB)
[30436.179846] sdb: sdb1 sdb2
...
```

- Unmount all SD card partitions.
- Launch GParted:

```
host:~$ sudo gparted
```





5.8.1.1 Expand rootfs


Warning

Running gparted on host systems which are using resize2fs version 1.46.6 and older (e.g. Ubuntu 22.04) are not able to expand the ext4 partition created with Yocto Mickledore and newer. This is due to a new default option in resize2fs which causes a incompatibility. See [release notes](#).


- Choose your SD card device at the drop-down menu on the top right
- Choose the ext4 root partition and click on resize:



- Drag the slider as far as you like or enter the size manually.



- Confirm your entry by clicking on the “Change size” button.




- To apply your changes, press the green tick.

- Now you can mount the root partition and copy e.g. the phytec-vision-image-imx8mp-phyflex-phyvip-1.wic image to it. Then unmount it again:


```
host:~$ sudo cp phytec-vision-image-imx8mp-phyflex-phyvip-1.wic /mnt/ ; sync
host:~$ umount /mnt
```

5.8.1.2 Create the Third Partition

- Choose your SD card device at the drop-down menu on the top right



- Choose the bigger unallocated area and press “New”:



Create new Partition

Minimum size: 1 MiB Maximum size: 2878 MiB

Free space preceding (MiB): Create as: Primary Partition


New size (MiB): Partition name:

Free space following (MiB): File system: ext4

Align to: MiB Label:

Cancel Add

- Click “Add”



/dev/sde - GParted

GParted Edit View Device Partition Help

/dev/sde (3.71 GiB)

/dev/sde2
848.81 MiB


New Partition #1
2.81 GiB

Partition	File System	Label	Size	Used	Unused	Flags
unallocated	unallocated		4.00 MiB	---	---	
/dev/sde1	fat16	boot	65.00 MiB	40.13 MiB	24.87 MiB	boot, lba
/dev/sde2	ext4	root	848.81 MiB	848.81 MiB	0.00 B	
New Partition #1	ext4		2.81 GiB	---	---	

Create Primary Partition #1 (ext4, 2.81 GiB) on /dev/sde

1 operation pending

- Confirm your changes by pressing the green tick.



- Now you can mount the new partition and copy e.g. phytec-vision-image-imx8mp-phyflex-phyvip-1.wic image to it. Then unmount it again:

```
host:~$ sudo mount /dev/sde3 /mnt
host:~$ sudo cp phytec-vision-image-imx8mp-phyflex-phyvip-1.wic /mnt/ ; sync
host:~$ umount /mnt
```

5.9 5.9 Ampliphy-boot

Ampliphy-boot is a collection of bootscripts for booting the Ampliphy distro, which work uniformly across different PHYTEC platforms. The bootscripts are developed to work with the [standard boot](#) U-Boot feature. More details about ampliphy-boot, the bootscripts and their Yocto implementation can be found in the Yocto Reference Manual (walnascar).

Ampliphy-boot currently provides these bootscripts:

```
mmc_boot
mmc_boot_fit
net_boot_fit
spi_boot_fit
```

Not all platforms support all scripts. The **phyVIP-i.MX 8M Plus Vision Integration Platform** supports the following bootscripts:

```
mmc_boot_fit
net_boot_fit
```

5.9.1 5.9.1 U-Boot standard boot

U-Boot standard boot is a built-in way to boot different operating systems. It handles scanning devices for the supported [boot methods](#) and starts the [bootflow](#).

For ampliphy-boot the `script` bootmethod is used. This means a bootscript is placed inside the boot partition on the bootdevice that should be booted. This bootscript contains the logic for booting the ampliphy distro that previously was stored inside the U-Boot environment. This leads to the environment being a lot less cluttered and better readable for the user.

U-Boot standard boot uses and sets some environment variables, which are documented [here](#). The most interesting variables for the BSP user are:

ip_dyn

determines if the bootscript is loaded with dhcp (dynamic ip) or tftp (static ip). Values can be [yes/no], [true/false], [1/0]. If the variable is not set, dhcp will be used as default.

boot_script_dhcp

Sets the name of the bootscript loaded from network.

boot_targets

Determines the scanned bootsources and their order.

bootmethods

Determines the scanned bootmethods and their order.

Default values for `boot_targets` and `bootmethods` for standard boot can be set in the devicetree or in the environment. The environment variables will overwrite the devicetree values if set.

For the **phyVIP-i.MX 8M Plus Vision Integration Platform**, the default values are defined in the U-Boot devicetree (e.g. `arch/arm/dts/imx8mp-phyflex-phyvip-u-boot.dtsi`):

```
bootstd {
    bootph-verify;
    compatible = "u-boot,boot-std";

    filename-prefixes = "/", "/boot/";
    bootdev-order = "mmc2", "mmc1", "ethernet";

    rauc {
        compatible = "u-boot,distro-rauc";
    };

    script {
        compatible = "u-boot,script";
    };
};
```

The `filename-prefixes` property describes the paths that will be searched for the bootscripts. In this case this is the root of the partition as well as the boot folder. The `bootdev-order` property sets the default value for the `boot_targets` variable. The supported `bootmethods` will also be named. In this case the `efi`, `rauc` and `script` bootmethods are supported.

5.10 5.10 Working with FIT images

Flattened Image Tree (FIT) images can be used to pack binaries into a devicetree blob. They are used in our BSPs to pack the kernel, devicetree and devicetree overlays.

5.10.1 5.10.1 Show FIT image content

Printing the FIT image content on the host can be done with the following commands:

```
host:~$ dumpimage -l fitImage
```

or

```
host:~$ mkimage -l fitImage
```

Alternatively the content can also be printed on target in U-Boot:

```
u-boot=> load mmc ${mmcdev}:1 ${loadaddr} fitImage
u-boot=> iminfo ${loadaddr}
```

5.10.2 5.10.2 Building FIT images standalone

For building a FIT, an image tree source file (*.its) and binaries to be packed are needed. The `fitImage.its` file and kernel and devicetree binaries can be fetched from the Yocto deploy folder, the BSP downloads page or can be built standalone. The paths of the binaries in the `data` properties may need adjustment. For the kernel, the default data property will mention `linux.bin`, but the binary from the Yocto build is named `Image`.

Default:

```
images {
    kernel-1 {
        description = "Linux kernel";
        data = /incbin/("linux.bin");
        type = "kernel";
    }
    ...
}
```

Customized:

```
images {
    kernel-1 {
        description = "Linux kernel";
        data = /incbin/("Image");
        type = "kernel";
    }
    ...
}
```

Data properties for devicetrees also may need to be adjusted. After that the FIT image can be built using the `mkimage` command.

```
mkimage -f fitImage.its fitImage
```

5.10.3 5.10.3 Adding overlays to a FIT image

Adding overlays to the FIT image can be done by editing the `fitImage.its` file. For that the overlays need to be added to the `images` node and configurations need to be created for them. Already existing overlays can be used for reference. Only the names need to be adjusted.

6 DEVICE TREE (DT)

6.1 6.1 Introduction

The following text briefly describes the Device Tree and can be found in the Linux kernel Documentation (<https://docs.kernel.org/devicetree/usage-model.html>)

“The “Open Firmware Device Tree”, or simply Devicetree (DT), is a data structure and language for describing hardware. More specifically, it is a description of hardware that is readable by an operating system so that the operating system doesn’t need to hard code details of the machine.”

The kernel documentation is a really good source for a DT introduction. An overview of the device tree data format can be found on the device tree usage page at devicetree.org.

6.2 6.2 PHYTEC i.MX 8M Plus BSP Device Tree Concept

The following sections explain some rules PHYTEC has defined on how to set up device trees for our i.MX 8M Plus SoC-based boards.

6.2.1 6.2.1 Device Tree Structure

- *Module.dtsi* - Module includes all devices mounted on the SoM, such as PMIC and RAM.
- *Board.dts* - include the module dtsi file. Devices that come from the i.MX 8M Plus SoC but are just routed down to the carrier board and used there are included in this dts.
- *Overlay.dts* - enable/disable features depending on optional hardware that may be mounted or missing on SoM or baseboard (e.g SPI flash or PEB-AV-10)

From the root directory of the Linux Kernel our devicetree files for i.MX 8 platforms can be found in `arch/arm64/boot/dts/freescale/`.

6.2.2 6.2.2 Device Tree Overlay

Device Tree overlays are device tree fragments that can be merged into a device tree during boot time. These are for example hardware descriptions of an expansion board. They are instead of being added to the device tree as an extra include, now applied as an overlay. They also may only contain setting the status of a node depending on if it is mounted or not. The device tree overlays are placed next to the other device tree files in our Linux kernel repository in the folder `arch/arm64/boot/dts/freescale/`.

Available overlays for `imx8mp-phyflex-phyvip-1.conf` are:

`imx8mp-phyflex-phyvip.dtb`
`imx8mp-phyflex-phyvip-peb-b-006.dtb`

(continues on next page)

(continued from previous page)

```

imx8mp-phyflex-phyvip-temperature.dtbo
imx8mp-phyflex-vm016-csi1.dtbo
imx8mp-phyflex-vm016-csi1-fpdlink-port0.dtbo
imx8mp-phyflex-vm016-csi1-fpdlink-port1.dtbo
imx8mp-phyflex-vm016-csi2.dtbo
imx8mp-phyflex-vm016-csi2-fpdlink-port0.dtbo
imx8mp-phyflex-vm016-csi2-fpdlink-port1.dtbo
imx8mp-phyflex-vm017-csi1.dtbo
imx8mp-phyflex-vm017-csi1-fpdlink-port0.dtbo
imx8mp-phyflex-vm017-csi1-fpdlink-port1.dtbo
imx8mp-phyflex-vm017-csi2.dtbo
imx8mp-phyflex-vm017-csi2-fpdlink-port0.dtbo
imx8mp-phyflex-vm017-csi2-fpdlink-port1.dtbo
imx8mp-phyflex-vm020-csi1.dtbo
imx8mp-phyflex-vm020-csi1-fpdlink-port0.dtbo
imx8mp-phyflex-vm020-csi1-fpdlink-port1.dtbo
imx8mp-phyflex-vm020-csi2.dtbo
imx8mp-phyflex-vm020-csi2-fpdlink-port0.dtbo
imx8mp-phyflex-vm020-csi2-fpdlink-port1.dtbo
imx8mp-phyflex-vm024-csi1.dtbo
imx8mp-phyflex-vm024-csi2.dtbo

```

Otherwise you can show the content of a FIT image including all overlay configs in the FIT image with this command in Linux:

```
host:~$ dumpimage -l fitImage
```

or in U-Boot:

```

u-boot=> load mmc ${mmcdev}:1 ${loadaddr} fitImage
u-boot=> iminfo

```

The usage of overlays can be configured during runtime in Linux or U-Boot. Overlays are applied during the boot process in the bootloader after the boot command is called and before the kernel is loaded. The next sections explain the configuration in more detail.

6.2.2.1 Loading overlays

6.2.2.1.1 FIT image bootscripts

In ampliphy-boot FIT image bootscripts, the `fit_overlay_conf` U-Boot environment variable contains a number-sign (#) separated list of overlay configurations that will be applied during boot. This variable is used for overlays describing expansion boards and cameras that can not be detected during run time. The overlays listed in the `fit_overlay_conf` variable must be included in the FIT image. Overlays set in the `$KERNEL_DEVICE_TREE` Yocto machine variable will automatically be added to the FIT image.

The `fit_overlay_conf` variable can either be set directly in the U-Boot environment or can be part of the external `overlays.txt` environment file. By default, the `fit_overlay_conf` variable comes from the external `overlays.txt` environment file which is located in the boot partition. Otherwise if the `fit_overlay_conf` variable is set in U-Boot environment, the `overlays.txt` file value will not be loaded. The content from the file is defined in the Yocto recipe `bootenv` found in `meta-phytec`: <https://github.com/phytec/meta-phytec/tree/walnascar/recipes-bsp/bootenv> You can read and write the file on booted target from linux:

```
target:~$ cat /boot/overlays.txt
fit_overlay_conf=conf-imx8mp-phyflex-phyvip-overlay1.dtbo#conf-imx8mp-phyflex-phyvip-overlay2.
↵dtbo
```

Changes will take effect after the next reboot. If no `overlays.txt` file is available the `fit_overlay_conf` variable can be set directly in the U-Boot environment.

```
u-boot=> setenv fit_overlay_conf conf-imx8mp-phyflex-phyvip-overlay1.dtbo
u-boot=> printenv fit_overlay_conf
fit_overlay_conf=conf-imx8mp-phyflex-phyvip-overlay1.dtbo
u-boot=> boot
```

6.2.2.1.2 Non FIT image bootscripts

In ampliphy-boot non FIT image bootscripts, the `overlays` U-Boot environment variable contains a space () separated list of overlays that will be applied during boot. The `overlays` variable can also be set directly in the U-Boot environment or can be part of the external `overlays.txt` environment file.

By default, the `overlays` variable comes from the external `overlays.txt` environment file which is located in the boot partition. Otherwise if the `overlays` variable is set in U-Boot environment, the `overlays.txt` file value will not be loaded. The content from the file is defined in the Yocto recipe `bootenv` found in meta-phytec: <https://github.com/phytec/meta-phytec/tree/walnascar/recipes-bsp/bootenv> You can read and write the file on booted target from linux:

```
target:~$ cat /boot/overlays.txt
overlays=imx8mp-phyflex-phyvip-overlay1.dtbo imx8mp-phyflex-phyvip-overlay2.dtbo
```

Changes will take effect after the next reboot. If no `overlays.txt` file is available the `overlays` variable can be set directly in the U-Boot environment.

```
u-boot=> setenv overlays imx8mp-phyflex-phyvip-overlay1.dtbo
u-boot=> printenv overlays
overlays=imx8mp-phyflex-phyvip-overlay1.dtbo
u-boot=> boot
```

6.2.3 6.2.3 Change U-boot Environment from Linux on Target

Libubootenv is a tool included in our images to modify the U-Boot environment of Linux on the target machine.

Print the U-Boot environment using the following command:

```
target:~$ fw_printenv
```

Modify a U-Boot environment variable using the following command:

```
target:~$ fw_setenv <variable> <value>
```

Caution

Libubootenv takes the environment selected in a configuration file. The environment to use is inserted there, and by default it is configured to use the eMMC environment (known as the default used environment).

If the eMMC is not flashed or the eMMC environment is deleted, libubootenv will not work. You should modify the `/etc/fw_env.config` file to match the environment source that you want to use.

6.2.4 6.2.4 PHYTEC i.MX 8M Plus FPSC Device Tree Differences

Due to the standardization of signals and interfaces that PHYTEC FPSC SoMs must supply, there is potential for confusion of the names as defined by the SoC manufacturer, in Software by the device tree and in the FPSC specification.

Example of i.MX 8M Plus SoC with a Libra i.MX 8M Plus FPSC board:

SoC vendor	device tree	FPSC Spec.
uSDHC1	usdhc1	SDIO
uSDHC2	usdhc2	SD-Card
uSDHC3	usdhc3	

To reduce confusion, interfaces in the device tree are annotated with the FPSC specification name.

Example (taken from `imx8mp-phycore-fpsc.dtsi`)

```
&usdhc2 { /* FPSC SDCARD */
    pinctrl-0 = <&pinctrl_usdhc2>;
    pinctrl-1 = <&pinctrl_usdhc2_100mhz>;
    pinctrl-2 = <&pinctrl_usdhc2_200mhz>;
    pinctrl-names = "default", "state_100mhz", "state_200mhz";
    sd-uhs-sdr104;
    vmmc-supply = <&reg_usdhc2_vmmc>;
    vqmmc-supply = <&ldo5>;
};

pinctrl_usdhc2: usdhc2grp {
    fsl,pins = <
        MX8MP_IOMUXC_SD2_CD_B_USDHC2_CD_B    0x40    /* SDCARD_CD */
        MX8MP_IOMUXC_SD2_WP_USDHC2_WP        0x40    /* SDCARD_WP */
        MX8MP_IOMUXC_SD2_CLK_USDHC2_CLK      0x190   /* SDCARD_CLK */
        MX8MP_IOMUXC_SD2_CMD_USDHC2_CMD      0x1d0   /* SDCARD_CMD */
        MX8MP_IOMUXC_SD2_DATA0_USDHC2_DATA0  0x1d0   /* SDCARD_DATA0 */
        MX8MP_IOMUXC_SD2_DATA1_USDHC2_DATA1  0x1d0   /* SDCARD_DATA1 */
        MX8MP_IOMUXC_SD2_DATA2_USDHC2_DATA2  0x1d0   /* SDCARD_DATA2 */
        MX8MP_IOMUXC_SD2_DATA3_USDHC2_DATA3  0x1d0   /* SDCARD_DATA3 */
        MX8MP_IOMUXC_GPI01_I004_USDHC2_VSELECT 0xc0
    >;
};
```

In the first node, the interface as defined per FPSC is annotated: `FPSC SDCARD`. In the second node, the signal names for this (SDCARD) interface are annotated: `SDCARD_WP`, etc. .

6.2.4.1 FPSC Device Tree Structure

In addition to the usual PHYTEC device tree structure, all FPSC defined signals are also predefined in the SoM device tree as can be seen in the example above. There is nothing connected to the SDCARD controller (on the SoM), yet the FPSC SoM must preconfigure SDCARD signals per the FPSC specification. The

result is that on the baseboard, no care needs to be taken in regards to signal line configuration (muxing); this is done already.

Example cont. (imx8mp-libra-rdk-fpsc.dts)

```
/* SD-Card */
&usdhc2 {
    assigned-clocks = <&clk IMX8MP_CLK_USDHC2>;
    assigned-clock-rates = <200000000>;
    bus-width = <4>;
    disable-wp;
    status = "okay";
};
```

6.2.4.2 Using the phyFLEX-i.MX 8M Plus FPSC Gamma for non-FPSC-conformant carrier boards

Even though signals for pins are predefined within the FPSC world, the phyFLEX-i.MX 8M Plus FPSC Gamma may also be used in custom, non-conforming (with regards to FPSC), i.e. “phyCORE”, design context.

Example:

SDCARD interface is not needed on custom board, but an additional i2c interface is required.

The following could be done

```
&i2c4 {
    pinctrl-0 = <&pinctrl_i2c4>;
    pinctrl-names = "default";
    status = "okay";
};

pinctrl_i2c4: i2c4grp {
    fsl,pins = <
        MX8MP_IOMUXC_SD2_DATA0__I2C4_SDA 0x400001c2
        MX8MP_IOMUXC_SD2_DATA1__I2C4_SCL 0x400001c2
    >;
};
```

Note that SD2_DATA0 pad, even though muxed in the phyFLEX-i.MX 8M Plus FPSC Gamma dtsi file, is not used since the interface controller line it is muxed to is not activated with a `status = "okay";` property.

7 ACCESSING PERIPHERALS

To find out which boards and modules are supported by the release of PHYTEC’s phyFLEX-i.MX 8M Plus FPSC Gamma BSP described herein, visit [our BSP](#) web page and click the corresponding BSP release in the download section. Here you can find all hardware supported in the columns “Hardware Article Number” and the correct machine name in the corresponding cell under “Machine Name”.

To achieve maximum software reuse, the Linux kernel offers a sophisticated infrastructure that layers software components into board-specific parts. The BSP tries to modularize the kit features as much as possible. When a customized baseboard or even a customer-specific module is developed, most of the software support can be reused without error-prone copy-and-paste. The kernel code corresponding to the boards can be found in device trees (DT) in the kernel repository under `arch/arm64/boot/dts/freescale/*.dts`.

In fact, software reuse is one of the most important features of the Linux kernel, especially of the ARM implementation which always has to fight with an insane number of possibilities of the System-on-Chip CPUs. The whole board-specific hardware is described in DTs and is not part of the kernel image itself. The hardware description is in its own separate binary, called the Device Tree Blob (DTB) (section [device tree](#)).

Please read section PHYTEC i.MX 8M Plus BSP Device Tree Concept to get an understanding of our i.MX 8 BSP device tree model.

The following sections provide an overview of the supported hardware components and their operating system drivers on the i.MX 8 platform. Further changes can be ported upon customer request.

7.1 7.1 i.MX 8M Plus Pin Muxing

The i.MX 8M Plus SoC contains many peripheral interfaces. In order to reduce package size and lower overall system cost while maintaining maximum functionality, many of the i.MX 8M Plus terminals can multiplex up to eight signal functions. Although there are many combinations of pin multiplexing that are possible, only a certain number of sets, called IO sets, are valid due to timing limitations. These valid IO sets were carefully chosen to provide many possible application scenarios for the user.

Please refer to our Hardware Manual or the NXP i.MX 8M Plus Reference Manual for more information about the specific pins and the muxing capabilities.

The IO set configuration, also called muxing, is done in the Device Tree. The driver `pinctrl-single` reads the DT’s node `fsl,pins`, and does the appropriate pin muxing.

The following is an example of the pin muxing of the UART3 device in `imx8mp-phyflex-fpsc-g-som.dtsi`:

```
pinctrl_uart4: uart4grp {
    fsl,pins = <
        MX8MP_IOMUXC_UART4_RXD__UART4_DCE_RX 0x140    /* UART3_RXD */
        MX8MP_IOMUXC_UART4_TXD__UART4_DCE_TX 0x140    /* UART3_TXD */
    >;
};
```

The first part of the string `MX8MP_IOMUXC_UART4_RXD__UART4_DCE_RX` names the pad (in this example `UART4_RX`). The second part of the string (`UART3_DCE_RX`) is the desired muxing option for this pad. The pad setting value (hex value on the right) defines different modes of the pad, for example, if internal pull resistors are activated or not. In this case, the internal resistors are disabled.

The device tree representation for UART3 pinmuxing: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-fpsc-g-som.dtsi#L612>

7.2 7.2 RS232/RS485

The FPSC Standard supports 3 UART units. On the phyVIP-i.MX 8M Plus, TTL level signals of UART3 (the standard console) are routed to a connector (*X8*).

Warning

The TTL level signals of (*X8*) are 1.8V compatible and not 3.3V tolerant. Please use a USB to TTL serial adapter with 1.8V logic level only!

UART2 interface is used for the onboard Bluetooth module and thus not available for external use.

UART1 interface is thus the only one brought out to phyVIP board connector (*CON1*) and can be used for general purpose RS232 or RS485 communication.

Note

With the phyVIP default expansion board PEB-B-006, UART1 is connected to a RS-485 transceiver chip capable of translating the UART1 signals to RS-485 differential signals available at the connector *X12*.

7.2.1 7.2.1 RS232

- Display the current settings of a terminal in a human-readable format:

```
target:~$ stty -a
```

- Configuration of the UART interface can be done with `stty`. For example:

```
target:~$ stty -F /dev/ttymxcl 115200 crtscts raw -echo
```

- With a simple echo and cat, basic communication can be tested. Example:

```
target:~$ echo 123 > /dev/ttymxcl
```

```
host:~$ cat /dev/ttyUSB2
```

The host should print out "123".

7.2.2 7.2.2 RS485

Hint

Remember to use bus termination resistors of 120 Ohm at each end of the bus, when using longer cables.

For easy testing, look at the linux-serial-test. This tool is called the IOCTL for RS485 and sends a constant stream of data.

```
target:~$ linux-serial-test -p /dev/ttymxcl -b 115200 --rs485 0
```

More information about the linux-serial-test tool and its parameters can be found here: [linux-serial-test](#)

The linux-serial-test will automatically set ioctls, but they can also be set manually with rs485conf.

You can show the current config with:

```
target:~$ rs485conf /dev/ttymxcl
```


You can show all options with:

```
target:~$ rs485conf /dev/ttymxcl -h
```

Documentation for calling the IOCTL within c-code is described in the Linux kernel documentation: <https://www.kernel.org/doc/Documentation/serial/serial-rs485.txt>

7.2.2.1 RS485 half-duplex

For half-duplex mode your connection setup should look like this:



Note

Note that TX+ corresponds to signal A and TX- corresponds to signal B.

Which function is on which pin is described in the hardware manual.

For half-duplex mode you can set the ioctls manually like this:

```
target:~$ rs485conf /dev/ttymxcl -e 1 -r 0
target:~$ rs485conf /dev/ttymxcl
= Current configuration:
RS485 enabled:           true
RTS on send:             high
RTS after send:          low
RTS delay before send:   0
RTS delay after send:    0
```

(continues on next page)

(continued from previous page)

```
Receive during sending data: false
Bus termination enabled: false
```

Then you can test if sending and receiving works like this:

```
target1:~$ cat /dev/ttymxcl
target2:~$ echo test > /dev/ttymxcl
```

You should see “test” printed out on target1. You can also switch the roles and send on target2 and receive on target1.

Alternatively you can also test with the linux-serial-test tool:


```
target1:~$ linux-serial-test -s -e -f -p /dev/ttymxcl -b 115200 --rs485 0 -t -i 8
...
/dev/ttymxcl: count for this session: rx=57330, tx=0, rx err=0
target2:~$ linux-serial-test -s -e -f -p /dev/ttymxcl -b 115200 --rs485 0 -r -o 5
...
/dev/ttymxcl: count for this session: rx=0, tx=57330, rx err=0
```

In this example target1 will be the receiver and target2 will be the transmitter. You should also be able to switch the roles. Remember to first start the receiver and then the transmitter immediately after. The receiver will receive for 8 sec and the transmitter will send for 5 sec. The receiver needs to receive for a bit longer than the transmitter sends. At the end the program will print the final “count for this session”. There you can check, that all transmitted frames were received.

All the tests are target to target, but can also be done with host to target with a USB to rs485 converter. You may need to adjust the interfaces then.

7.2.2.2 RS485 full-duplex

For full-duplex mode your connection setup should look like this:



Which function is on which pin is described in the hardware manual.

For full-duplex mode you can set the ioctls manually like this:

```
target:~$ rs485conf /dev/ttymxcl -e 1 -r 1
target:~$ rs485conf /dev/ttymxcl
= Current configuration:
RS485 enabled:           true
```

(continues on next page)

(continued from previous page)

```
RTS on send:           high
RTS after send:        low
RTS delay before send: 0
RTS delay after send:  0
Receive during sending data: true
Bus termination enabled: false
```

Also here you can do the echo test to see if sending and receiving works:

```
target1:~$ cat /dev/ttymxcl
target2:~$ echo test > /dev/ttymxcl
```

You should see “test” printed out on target1. You can also switch the roles and send on target2 and receive on target1.

To check if the full-duplex operation works, you need to use the linux-serial-test tool:

```
target1:~$ linux-serial-test -s -e -f -p /dev/ttymxcl -b 115200 --rs485 0 -o 10 -i 15 -W 2
...
/dev/ttymxcl: count for this session: rx=114660, tx=118755, rx err=0
target2:~$ linux-serial-test -s -e -f -p /dev/ttymxcl -b 115200 --rs485 0 -o 10 -i 15 -W 2
...
/dev/ttymxcl: count for this session: rx=118755, tx=114660, rx err=0
```

In this example both targets will send and receive simultaneously. They will receive for 15sec and send for 10sec. The receiver needs to receive a bit longer, so that all sent messages will get received. Remember to start both targets almost simultaneously. A small difference in start time is accounted for with the `-W 2` option. At the end the program will print the final “count for this session”. There you can check that all transmitted frames were received.

All the test examples are target to target, but can also be done with host to target with a USB to rs485 converter. You may need to adjust the interfaces for commands to work on the host then.

The device tree representation for RS485 with default expansion board PEB-B-006: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip-peb-b-006.dts#L170>

7.3 Ethernet

phyVIP-i.MX 8M Plus provides two Ethernet interfaces. However, with the default expansion board PEB-B-006, only one gigabit Ethernet interface is available.

All interfaces offer a standard Linux network port that can be programmed using the BSD socket interface. The whole network configuration is handled by the systemd-networkd daemon. The relevant configuration files can be found on the target in `/lib/systemd/network/` as well as the BSP in `meta-ampliphy/recipes-core/systemd/systemd-conf`.

IP addresses can be configured within `*.network` files. The interfaces are configured to static IP as default. The default IP address and netmask for end1 is:

```
end1: 192.168.3.11/24
```

To configure end1 to dynamic IP over DHCP, go to `/lib/systemd/network/*-end1.network` and delete the line:

```
Address=192.168.3.11/24
```

The DT Ethernet setup might be split into two files depending on your hardware configuration: the module DT and the board-specific DT. The device tree set up for the ethernet where the PHY is populated on the SoM can be found here: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-fpsc-g-som.dtsi#L94>.

7.3.1 7.3.1 Network Environment Customization

7.3.1.1 U-boot network-environment

- To find the Ethernet settings in the target bootloader:

```
u-boot=> printenv ipaddr serverip netmask
```

- With your development host set to IP 192.168.3.10 and netmask 255.255.255.0, the target should return:

```
u-boot=> printenv ipaddr serverip netmask
ipaddr=192.168.3.11
serverip=192.168.3.10
netmask=255.225.255.0
```

- If you need to make any changes:

```
u-boot=> setenv <parameter> <value>
```

<parameter> should be one of ipaddr, netmask, gatewayip or serverip. <value> will be the actual value of the chosen parameter.

- The changes you made are temporary for now. To save these:

```
u-boot=> saveenv
```

Here you can also change the IP address to DHCP instead of using a static one.

- Configure:

```
u-boot=> setenv ip dhcp
```

- Set up paths for TFTP and NFS. A modification could look like this:

```
u-boot=> setenv nfsroot /home/user/nfssrc
```

Please note that these modifications will only affect the bootloader settings.

Tip

Variables like nfsroot and netargs can be overwritten by the U-boot External Environment. For netboot, the external environment will be loaded from tftp. For example, to locally set the nfsroot variable in a bootenv.txt file, locate the tftproot directory:

```
nfsroot=/home/user/nfssrc
```

There is no need to store the info on the target. Please note that this does not work for variables like `ipaddr` or `serveraddr` as they need to be already set when the external environment is being loaded.

7.3.1.1.1 Multiple network interfaces in U-Boot

Some boards may have support for multiple network interfaces in U-Boot. U-Boot lists all supported network interfaces during startup. This generic example illustrates support for multiple network interfaces:

```
Net:
eth0: ethernet@10000000, eth1: ethernet@10010000
```

There are three environment variables that will contribute to the selection of the ethernet interface for example when loading a file over tftp.

ethrotate

Determines if other interfaces will be probed if the first one fails. Valid values are yes/no. Default is yes if unset.

ethact

Sets the current active ethernet interface. This interface will be used (first). It is usually unset at boot, but will be set automatically by U-Boot to the last used interface as soon as ethernet is used for the first time. U-Boot will overwrite existing values. Only one ethernet interface can be set, lists are not allowed.

ethprime

This will determine the preferred ethernet interface if `ethact` is unset. Otherwise this will be ignored. Only one ethernet interface can be set, lists are not allowed.

Note

The variables `ipaddr`, `serverip` and `netmask` are global and not specific to a certain ethernet interface, so they might need to be adapted to the ethernet interface in use.

7.3.1.2 Kernel network-environment

- Find the ethernet settings for `end1` in the target kernel:

```
target:~$ ip -statistics address show end1
2: end1: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state UP group default qlen
↪ 1000
    link/ether 50:2d:f4:19:d6:33 brd ff:ff:ff:ff:ff:ff
    RX:  bytes  packets  errors  dropped  missed  mcast
         0         0         0         0         0         0
    TX:  bytes  packets  errors  dropped  carrier  collsns
         0         0         0         0         0         0
```

- Temporary adaption of the `end1` configuration:

```
target:~$ ip address add 192.168.3.11/24 dev end1
```

7.4 7.4 WLAN/Bluetooth

For WLAN and Bluetooth support, phyVIP utilizes the BDE-3P-BW335X module from BDE Technology Inc. This modules support 2.4-GHz & 5-GHz Dual-Band Wi-Fi 6 and Bluetooth Low Energy. The combo module is based around Texas Instruments CC3351 chip.

More information about the module can be found at <https://www.bdecomm.com/uploads/2026429/20260429101391389138.pdf>

Note

Bluetooth functionality is not yet available due to missing drivers. Please contact PHYTEC support for more information on the driver availability.

7.4.1 7.4.1 Connecting to a WLAN Network

First set the correct regulatory domain for your country:

```
target:~$ iw reg set DE
target:~$ iw reg get
```

You will see:

```
country DE: DFS-ETSI
(2400 - 2483 @ 40), (N/A, 20), (N/A)
(5150 - 5250 @ 80), (N/A, 20), (N/A), NO-OUTDOOR
(5250 - 5350 @ 80), (N/A, 20), (0 ms), NO-OUTDOOR, DFS
(5470 - 5725 @ 160), (N/A, 26), (0 ms), DFS
(57000 - 66000 @ 2160), (N/A, 40), (N/A)
```

Set up the wireless interface:

```
target:~$ ip link
target:~$ ip link set up dev wlan0
```

Now you can scan for available networks:

```
target:~$ iw wlan0 scan | grep SSID
```

You can use a cross-platform supplicant with support for WEP, WPA, and WPA2 called wpa_supplicant for an encrypted connection.

To do so, add the network-credentials to the file /etc/wpa_supplicant.conf:

```
country=DE
network={
    ssid="<SSID>"
    proto=WPA2
    psk="<KEY>"
}
```

Now a connection can be established:

```
target:~$ wpa_supplicant -D nl80211 -c /etc/wpa_supplicant.conf -i wlan0 -B
```

This should result in the following output:

```
Successfully initialized wpa_supplicant
```

The ip address is automatically configured over DHCP. For other possible IP configurations, see section *Changing the Network Configuration* in the Yocto Reference Manual (walnascar).

7.5 7.5 SD card

The i.MX 8M Plus supports a slot for Secure Digital cards to be used as general-purpose block devices. These devices can be used in the same way as any other block device.

Warning

These kinds of devices are hot-pluggable. Nevertheless, you must ensure not to unplug the device while it is still mounted. This may result in data loss!

After inserting an SD card, the kernel will generate new device nodes in `/dev`. The full device can be reached via its `/dev/mmcblk1` device node. SD card partitions will show up as:

```
/dev/mmcblk1p<Y>
```

`<Y>` counts as the partition number starting from 1 to the max count of partitions on this device. The partitions can be formatted with any kind of file system and also handled in a standard manner, e.g. the mount and umount command work as expected.

Tip

These partition device nodes will only be available if the card contains a valid partition table ("hard disk" like handling). If no partition table is present, the whole device can be used as a file system ("floppy" like handling). In this case, `/dev/mmcblk1` must be used for formatting and mounting. The cards are always mounted as being writable.

DT configuration for the MMC (SD card slot) interface can be found here: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-fpsc-g-som.dtsi#L821> and <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip.dts#L187>

DT configuration for the e.MMC interface can be found here: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-fpsc-g-som.dtsi#L831>

7.6 7.6 e.MMC Devices

PHYTEC modules like phyFLEX-i.MX 8M Plus FPSC Gamma are populated with an e.MMC memory chip as the main storage. e.MMC devices contain raw Multi-Level Cells (MLC) or Triple-Level Cells (TLC) combined with a memory controller that handles ECC and wear leveling. They are connected via an SD/MMC interface to the i.MX 8M Plus and are represented as block devices in the Linux kernel like SD cards, flash drives, or hard disks.

The electric and protocol specifications are provided by JEDEC (<https://www.jedec.org/standards-documents/technology-focus-areas/flash-memory-ssds-ufs-embedded/e-mmc>). The e.MMC manu-

facturer's datasheet is relatively short and meant to be read together with the supported version of the JEDEC e.MMC standard.

PHYTEC currently utilizes the e.MMC chips with JEDEC Version 5.0 and 5.1

7.6.1 7.6.1 Extended CSD Register

e.MMC devices have an extensive amount of extra information and settings that are available via the Extended CSD registers. For a detailed list of the registers, see manufacturer datasheets and the JEDEC standard.

In the Linux user space, you can query the registers:

```
target:~$ mmc extcsd read /dev/mmcblk2
```

You will see:

```
=====
Extended CSD rev 1.7 (MMC 5.0)
=====

Card Supported Command sets [S_CMD_SET: 0x01]
[...]
```

7.6.2 7.6.2 Enabling Background Operations (BKOPS)

In contrast to raw NAND Flash, an e.MMC device contains a Flash Transfer Layer (FTL) that handles the wear leveling, block management, and ECC of the raw MLC or TLC. This requires some maintenance tasks (for example erasing unused blocks) that are performed regularly. These tasks are called **Background Operations (BKOPS)**.

By default (depending on the chip), the background operations may or may not be executed periodically, impacting the worst-case read and write latency.

The JEDEC Standard has specified a method since version v4.41 that the host can issue BKOPS manually. See the JEDEC Standard chapter Background Operations and the description of registers BKOPS_EN (Reg: 163) and BKOPS_START (Reg: 164) in the e.MMC datasheet for more details.

Meaning of Register BKOPS_EN (Reg: 163) Bit MANUAL_EN (Bit 0):

- Value 0: The host does not support the manual trigger of BKOPS. Device write performance suffers.
- Value 1: The host does support the manual trigger of BKOPS. It will issue BKOPS from time to time when it does not need the device.

The mechanism to issue background operations has been implemented in the Linux kernel since v3.7. You only have to enable BKOPS_EN on the e.MMC device (see below for details).

The JEDEC standard v5.1 introduces a new automatic BKOPS feature. It frees the host to trigger the background operations regularly because the device starts BKOPS itself when it is idle (see the description of bit AUTO_EN in register BKOPS_EN (Reg: 163)).

- To check whether *BKOPS_EN* is set, execute:

```
target:~$ mmc extcsd read /dev/mmcblk2 | grep BKOPS_EN
```

The output will be, for example:

```
Enable background operations handshake [BKOPS_EN]: 0x01
#0R
Enable background operations handshake [BKOPS_EN]: 0x00
```

Where value 0x00 means BKOPS_EN is disabled and device write performance suffers. Where value 0x01 means BKOPS_EN is enabled and the host will issue background operations from time to time.

- Enabling can be done with this command:

```
target:~$ target:~$ mmc --help

[...]
mmc bkops_en <auto|manual> <device>
    Enable the eMMC BKOPS feature on <device>.
    The auto (AUTO_EN) setting is only supported on eMMC 5.0 or newer.
    Setting auto won't have any effect if manual is set.
    NOTE! Setting manual (MANUAL_EN) is one-time programmable (unreversible) change.
```

- To set the BKOPS_EN bit, execute:

```
target:~$ mmc bkops_en manual /dev/mmcblk2
```

- To ensure that the new setting is taken over and the kernel triggers BKOPS by itself, shut down the system:

```
target:~$ poweroff
```

Tip

The BKOPS_EN bit is one-time programmable only. It cannot be reversed.

7.6.3 7.6.3 Reliable Write

There are two different Reliable Write options:

1. Reliable Write option for a whole e.MMC device/partition.
2. Reliable Write for single write transactions.

Tip

Do not confuse e.MMC partitions with partitions of a DOS, MBR, or GPT partition table (see the previous section).

The first Reliable Write option is mostly already enabled on the e.MMCs mounted on the phyFLEX-i.MX 8M Plus FPSC Gamma SoMs. To check this on the running target:

```
target:~$ mmc extcsd read /dev/mmcblk2 | grep -A 5 WR_REL_SET
Write reliability setting register [WR_REL_SET]: 0x1f
user area: the device protects existing data if a power failure occurs during a write o
peration
partition 1: the device protects existing data if a power failure occurs during a write
```

(continues on next page)

(continued from previous page)

```

operation
partition 2: the device protects existing data if a power failure occurs during a write
operation
partition 3: the device protects existing data if a power failure occurs during a write
operation
partition 4: the device protects existing data if a power failure occurs during a write
operation
--
Device supports writing EXT_CSD_WR_REL_SET
Device supports the enhanced def. of reliable write

```

Otherwise, it can be enabled with the mmc tool:

```

target:~$ mmc --help

[...]
mmc write_reliability set <-y|-n|-c> <partition> <device>
    Enable write reliability per partition for the <device>.
    Dry-run only unless -y or -c is passed.
    Use -c if more partitioning settings are still to come.
    NOTE! This is a one-time programmable (unreversible) change.

```

The second Reliable Write option is the configuration bit Reliable Write Request parameter (bit 31) in command CMD23. It has been used in the kernel since v3.0 by file systems, e.g. ext4 for the journal and user space applications such as fdisk for the partition table. In the Linux kernel source code, it is handled via the flag REQ_META.

Conclusion: ext4 file system with mount option data=journal should be safe against power cuts. The file system check can recover the file system after a power failure, but data that was written just before the power cut may be lost. In any case, a consistent state of the file system can be recovered. To ensure data consistency for the files of an application, the system functions fdatasync or fsync should be used in the application.

7.6.4 7.6.4 Resizing ext4 Root Filesystem

When flashing the SD card image to eMMC the ext4 root partition is not extended to the end of the eMMC. parted can be used to expand the root partition. The example works for any block device such as eMMC, SD card, or hard disk.

- Get the current device size:

```
target:~$ parted /dev/mmcblk2 print
```

- The output looks like this:

```

Model: MMC Q2J55L (sd/mmc)
Disk /dev/mmcblk2: 7617MB
Sect[ 1799.850385] mmcblk2: p1 p2
or size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number  Start   End     Size    Type    File system  Flags

```

(continues on next page)

(continued from previous page)

1	4194kB	72.4MB	68.2MB	primary	fat16	boot, lba
2	72.4MB	537MB	465MB	primary	ext4	

- Use parted to resize the root partition to the max size of the device:

```
target:~$ parted /dev/mmcblk2 resizepart 2 100%
Information: You may need to update /etc/fstab.

target:~$ parted /dev/mmcblk2 print
Model: MMC Q2J55L (sd/mmc)
Disk /dev/mmcblk2: 7617MB
Sector size (logical/physical): 512[ 1974.191657]   mmcblk2: p1 p2
B/512B
Partition Table: msdos
Disk Flags:

Number  Start   End     Size    Type    File system  Flags
  1      4194kB  72.4MB  68.2MB  primary fat16        boot, lba
  2      72.4MB  7617MB  7545MB  primary ext4
```

- Resize the filesystem to a new partition size:

```
target:~$ resize2fs /dev/mmcblk2p2
resize2fs 1.46.1 (9-Feb-2021)
Filesystem at /dev/mmcblk2p2 is mounted on /; on-line resizing required
[ 131.609512] EXT4-fs (mmcblk2p2): resizing filesystem
from 454136 to 7367680 blocks
old_desc_blocks = 4, new_desc_blocks = 57
[ 131.970278] EXT4-fs (mmcblk2p2): resized filesystem to 7367680
The filesystem on /dev/mmcblk2p2 is now 7367680 (1k) blocks long
```

Increasing the filesystem size can be done while it is mounted. But you can also boot the board from an SD card and then resize the file system on the eMMC partition while it is not mounted.

7.6.5 7.6.5 Enable pseudo-SLC Mode

eMMC devices use MLC or TLC (https://en.wikipedia.org/wiki/Multi-level_cell) to store the data. Compared with SLC used in NAND Flash, MLC or TLC have lower reliability and a higher error rate at lower costs.

If you prefer reliability over storage capacity, you can enable the pseudo-SLC mode or SLC mode. The method used here employs the enhanced attribute, described in the JEDEC standard, which can be set for continuous regions of the device. The JEDEC standard does not specify the implementation details and the guarantees of the enhanced attribute. This is left to the chipmaker. For the Micron chips, the enhanced attribute increases the reliability but also halves the capacity.

Warning

When enabling the enhanced attribute on the device, all data will be lost.

The following sequence shows how to enable the enhanced attribute.

- First obtain the current size of the eMMC device with:

```
target:~$ parted -m /dev/mmcblk2 unit B print
```

You will receive:

```
BYT;  
/dev/mmcblk2:63652757504B:sd/mmc:512:512:unknown:MMC S0J58X;;
```

As you can see this device has 63652757504 Byte = 60704 MiB.

- To get the maximum size of the device after pseudo-SLC is enabled use:

```
target:~$ mmc extcsd read /dev/mmcblk2 | grep ENH_SIZE_MULT -A 1
```

which shows, for example:

```
Max Enhanced Area Size [MAX_ENH_SIZE_MULT]: 0x000764  
i.e. 3719168 KiB  
--  
Enhanced User Data Area Size [ENH_SIZE_MULT]: 0x000000  
i.e. 0 KiB
```

Here the maximum size is 3719168 KiB = 3632 MiB.

- Now, you can set enhanced attribute for the whole device, e.g. 3719168 KiB, by typing:

```
target:~$ mmc enh_area set -y 0 3719168 /dev/mmcblk2
```

You will get:

```
Done setting ENH_USR area on /dev/mmcblk2  
setting OTP PARTITION_SETTING_COMPLETED!  
Setting OTP PARTITION_SETTING_COMPLETED on /dev/mmcblk2 SUCCESS  
Device power cycle needed for settings to take effect.  
Confirm that PARTITION_SETTING_COMPLETED bit is set using 'extcsd read' after power cycle
```

- To ensure that the new setting has taken over, shut down the system:

```
target:~$ poweroff
```

and perform a power cycle. It is recommended that you verify the settings now.

- First, check the value of ENH_SIZE_MULT which must be 3719168 KiB:

```
target:~$ mmc extcsd read /dev/mmcblk2 | grep ENH_SIZE_MULT -A 1
```

You should receive:

```
Max Enhanced Area Size [MAX_ENH_SIZE_MULT]: 0x000764  
i.e. 3719168 KiB  
--  
Enhanced User Data Area Size [ENH_SIZE_MULT]: 0x000764  
i.e. 3719168 KiB
```

- Finally, check the size of the device:

```
target:~$ parted -m /dev/mmcblk2 unit B print
BYT;
/dev/mmcblk2:31742492672B:sd/mmc:512:512:unknown:MMC S0J58X;
```

7.6.6 7.6.6 Erasing the Device

It is possible to erase the eMMC device directly rather than overwriting it with zeros. The eMMC block management algorithm will erase the underlying MLC or TLC or mark these blocks as discard. The data on the device is lost and will be read back as zeros.

- After booting from SD card execute:

```
target:~$ blkdiscard -f --secure /dev/mmcblk2
```

The option `--secure` ensures that the command waits until the eMMC device has erased all blocks. The `-f` (force) option disables all checking before erasing and it is needed when the eMMC device contains existing partitions with data.

Tip

```
target:~$ dd if=/dev/zero of=/dev/mmcblk2 conv=fsync
```

also destroys all information on the device, but this command is bad for wear leveling and takes much longer!

7.6.7 7.6.7 eMMC Boot Partitions

An eMMC device contains four different hardware partitions: user, boot1, boot2, and rpmb.

The user partition is called the User Data Area in the JEDEC standard and is the main storage partition. The partitions boot1 and boot2 can be used to host the bootloader and are more reliable. Which partition the i.MX 8M Plus uses to load the bootloader is controlled by the boot configuration of the eMMC device. The partition rpmb is a small partition and can only be accessed via a trusted mechanism.

Furthermore, the user partition can be divided into four user-defined General Purpose Area Partitions. An explanation of this feature exceeds the scope of this document. For further information, see the JEDEC Standard Chapter 7.2 Partition Management.

Tip

Do not confuse eMMC partitions with partitions of a DOS, MBR, or GPT partition table.

The current PHYTEC BSP does not use the extra partitioning feature of eMMC devices. The U-Boot is flashed at the beginning of the user partition. The U-Boot environment is placed at a fixed location after the U-Boot. An MBR partition table is used to create two partitions, a FAT32 boot, and ext4 rootfs partition. They are located right after the U-Boot and the U-Boot environment. The FAT32 boot partition contains the kernel and device tree.

With eMMC flash storage it is possible to use the dedicated boot partitions for redundantly storing the bootloader. The Bootloader environment still resides in the user area before the first partition. The user area also still contains the bootloader which the image first shipped during its initialization process. Below is an example, to flash the bootloader to one of the two boot partitions and switch the boot device via userspace commands.

7.6.7.1 Via userspace Commands

On the host, run:

```
host:~$ scp <bootloader> root@192.168.3.11:/tmp/
```

The partitions boot1 and boot2 are read-only by default. To write to them from user space, you have to disable `force_ro` in the sysfs.

To manually write the bootloader to the eMMC boot partitions, first disable the write protection:

```
target:~$ echo 0 > /sys/block/mmcblk2boot0/force_ro
target:~$ echo 0 > /sys/block/mmcblk2boot1/force_ro
```

Write the bootloader to the eMMC boot partitions:

```
target:~$ dd if=/tmp/<bootloader> of=/dev/mmcblk2boot0
target:~$ dd if=/tmp/<bootloader> of=/dev/mmcblk2boot1
```

The following table is for the offset of the i.MX 8M Plus SoC:

SoC	Offset User Area	Offset Boot Partition	e.MMC Device
i.MX 8M Plus	32 kiB	0 kiB	/dev/mmcblk2

After that set the boot partition from user space using the `mmc` tool:

(for 'boot0') :

```
target:~$ mmc bootpart enable 1 0 /dev/mmcblk2
```

(for 'boot1') :

```
target:~$ mmc bootpart enable 2 0 /dev/mmcblk2
```

To disable booting from the eMMC boot partitions simply enter the following command:

```
target:~$ mmc bootpart enable 0 0 /dev/mmcblk2
```

To explicitly enable booting from the eMMC user area, run:

```
target:~$ mmc bootpart enable 7 0 /dev/mmcblk2
```

7.6.7.2 Automatic failover

The ROM loader implements an automatic failover mechanism for eMMC boot partitions. If booting from the primary partition fails, the system automatically attempts to boot from the secondary partition. This failover is indicated by a change in the boot message from **Boot Stage: Primary boot** to **Boot Stage: Secondary boot**. This functionality is limited to boot0 and boot1 partitions and does not apply to the user area.

7.7 7.7 SPI Master

The i.MX 8M Plus controller has a FlexSPI and an ECSPI IP core included. The FlexSPI host controller supports two SPI channels with up to 4 devices. Each channel supports Single/Dual/Quad/Octal mode

data transfer (1/2/4/8 bidirectional data lines). The ECSPI controller supports 3 SPI interfaces with one dedicated chip selected for each interface. As chip selects should be realized with GPIOs, more than one device on each channel is possible.

Note

With the phyVIP-i.MX 8M Plus default expansion board PEB-B-006, the ECSPI interfaces are not brought out to any connector and thus not available externally.

The ECSPI2 interface is used for the onboard TPM2.0 chip. The definition of the SPI master node in the device tree can be found here:

<https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip.dts#L49>

7.8 7.8 GPIOs

The phyVIP has a set of pins especially dedicated to user I/Os. Those pins are connected directly to i.MX 8M Plus pins and are muxed as GPIOs. They are directly usable in Linux userspace. The processor has organized its GPIOs into five banks of 32 GPIOs each (GPIO1 – GPIO5) GPIOs. gpiochip0, gpiochip32, gpiochip64, gpiochip96, and gpiochip128 are the sysfs representation of these internal i.MX 8M Plus GPIO banks GPIO1 – GPIO5.

The GPIOs are identified as GPIO<X>_<Y> (e.g. GPIO5_07). <X> identifies the GPIO bank and counts from 1 to 5, while <Y> stands for the GPIO within the bank. <Y> is being counted from 0 to 31 (32 GPIOs on each bank).

By contrast, the Linux kernel uses a single integer to enumerate all available GPIOs in the system. The formula to calculate the right number is:

Linux GPIO number: <N> = (<X> - 1) * 32 + <Y>

Accessing GPIOs from userspace will be done using the libgpiod. It provides a library and tools for interacting with the Linux GPIO character device. Examples of some usages of various tools:

- Detecting the gpiochips on the chip:

```
target:~$ gpiodetect
gpiochip0 [30200000.gpio] (32 lines)
gpiochip1 [30210000.gpio] (32 lines)
gpiochip2 [30220000.gpio] (32 lines)
gpiochip3 [30230000.gpio] (32 lines)
gpiochip4 [30240000.gpio] (32 lines)
```

- Show detailed information about the gpiochips. Like their names, consumers, direction, active state, and additional flags:

```
target:~$ gpioinfo -c gpiochip0
```

- Read the value of a GPIO (e.g GPIO 20 from chip0):

```
target:~$ gpioget -c gpiochip0 20
```

- Set the value of GPIO 20 on chip0 to 0 and exit tool:

```
target:~$ gpioset -z -c gpiochip0 20=0
```

- Help text of gpioset shows possible options:

```
target:~$ gpioset --help
Usage: gpioset [OPTIONS] <line=value>...

Set values of GPIO lines.

Lines are specified by name, or optionally by offset if the chip option
is provided.
Values may be '1' or '0', or equivalently 'active'/'inactive' or 'on'/'off'.

The line output state is maintained until the process exits, but after that
is not guaranteed.

Options:
  --banner           display a banner on successful startup
  -b, --bias <bias>  specify the line bias
                     Possible values: 'pull-down', 'pull-up', 'disabled'.
                     (default is to leave bias unchanged)
  --by-name          treat lines as names even if they would parse as an offset
  -c, --chip <chip>  restrict scope to a particular chip
  -C, --consumer <name> consumer name applied to requested lines (default is 'gpioset')
  -d, --drive <drive> specify the line drive mode
                     Possible values: 'push-pull', 'open-drain', 'open-source'.
                     (default is 'push-pull')
  -h, --help          display this help and exit
  -l, --active-low    treat the line as active low
  -p, --hold-period <period>
                     the minimum time period to hold lines at the requested values
  -s, --strict        abort if requested line names are not unique
  -t, --toggle <period>[,period]...
                     toggle the line(s) after the specified period(s)
                     If the last period is non-zero then the sequence repeats.
  --unquoted         don't quote line names
  -v, --version        output version information and exit
  -z, --daemonize     set values then detach from the controlling terminal

Chips:
  A GPIO chip may be identified by number, name, or path.
  e.g. '0', 'gpiochip0', and '/dev/gpiochip0' all refer to the same chip.

Periods:
  Periods are taken as milliseconds unless units are specified. e.g. 10us.
  Supported units are 's', 'ms', and 'us'.

*Note*
  The state of a GPIO line controlled over the character device reverts to default
  when the last process referencing the file descriptor representing the device file
  exits.
  This means that it's wrong to run gpioset, have it exit and expect the line to continue
  being driven high or low. It may happen if given pin is floating but it must be
```

(continues on next page)

(continued from previous page)

↪ interpreted
as undefined behavior.

Warning

Some of the user IOs are used for special functions. Before using a user IO, refer to the schematic or the hardware manual of your board to ensure that it is not already in use.

7.8.1 7.8.1 GPIOs via sysfs

Warning

Accessing gpios via sysfs is deprecated and we encourage to use libgpiod instead.

Support to access GPIOs via sysfs is not enabled by default any more. It is only possible with manually enabling `CONFIG_GPIO_SYSFS` in the kernel configuration. To make `CONFIG_GPIO_SYSFS` visible in menuconfig the option `CONFIG_EXPERT` has to be enabled first.

You can also add this option for example to the defconfig you use in `arch/arm64/configs/` in the linux kernel sources. For our NXP based releases, this could be for example `imx8_phytec_defconfig`:

```
..
CONFIG_EXPERT=y
CONFIG_GPIO_SYSFS=y
..
```

Otherwise you can create a new config fragment. This is described in our Yocto Reference Manual.

7.9 7.9 I²C Bus

The i.MX 8M Plus contains several Multimaster fast-mode I²C modules. PHYTEC boards provide plenty of different I²C devices connected to the I²C modules of the i.MX 8M Plus. This section describes the basic device usage and its DT representation of some I²C devices integrated into our phyVIP.

The device tree node for i2c contains settings such as clock-frequency to set the bus frequency and the pin control settings including `scl-gpios` and `sda-gpios` which are alternate pin configurations used for bus recovery.

General I²C bus configuration from SoM (e.g. `imx8mp-phyflex-fpsc-g-som.dtsi`): <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-fpsc-g-som.dtsi#L205>

General I²C bus configuration from carrierboard (e.g. `imx8mp-phyflex-phyvip.dts`) <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip.dts#L63>

7.10 7.10 EEPROM

The system features three I²C EEPROM devices distributed across the SoM and carrier board:

On the phyFLEX-i.MX 8M Plus FPSC Gamma SoM:

- SoM Detection EEPROM (write-protected)
 - Bus: I2C-1
 - Address: 0x51
 - Purpose: Factory configuration for SoM identification
- User EEPROM
 - Bus: I2C-1
 - Address: 0x50
 - Purpose: Available for user applications

Device Tree Reference for SoM EEPROMs: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-fpsc-g-som.dtsi#L293>

And on the phyVIP carrier board:

- Board Detection EEPROM
 - Bus: I2C-2
 - Address: 0x51
 - Purpose: Reserved for carrier board identification

Device Tree Reference for Carrier Board: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip.dtsi#L67>

7.10.1 7.10.1 EEPROM reading and writing in U-Boot

In U-Boot the `i2c` command can be used for EEPROM read and write operations.

First set the correct `i2c` bus, where the eeprom is connected. `<bus-no>` is the `i2c` bus number of the EEPROM.

```
u-boot=> i2c dev <bus-no>
```

To read and print the first 32 bytes from EEPROM execute. `<addr>` is the `i2c` address of the EEPROM.

```
u-boot=> i2c md <addr> 0x00 0x20
```

To read the first 32 bytes from EEPROM into memory (`loadaddr`) execute:

```
u-boot=> i2c read <addr> 0x00 0x20 $loadaddr
```

To write `0xff` to the first 32 bytes execute:

```
u-boot=> i2c mw <addr> 0x00 0xff 0x20
```

To write 32 bytes to EEPROM from memory (`loadaddr`) execute:

```
u-boot=> i2c write $loadaddr <addr> 0x00 0x20
```

7.10.2 7.10.2 EEPROM reading and writing in Linux

Reading and writing can also be done via the `sysfs` in Linux. For this, find the correct path in `sysfs` first. It follows this logic: `/sys/class/i2c-dev/i2c-<bus-no>/device/<bus-no>-<addr>/eeprom` `<bus-no>` is the bus number of the EEPROM `<addr>` is the address of the EEPROM in 4 digits hex.

To read and print the EEPROM content as a hex number, execute:

```
target:~$ hexdump -C <eeprom-sysfs-path>
```

or:

```
target:~$ dd if=<eeprom-sysfs-path> | od -x
```

To fill the EEPROM with zeros use:

```
target:~$ dd if=/dev/zero of=<eeprom-sysfs-path>
```

7.11 7.11 RTC

RTCs can be accessed via `/dev/rtc*`. Because PHYTEC boards have often more than one RTC, there might be more than one RTC device file.

- To find the name of the RTC device, you can read its sysfs entry with:

```
target:~$ cat /sys/class/rtc/rtc*/name
```

- You will get, for example:

```
rtc-rv3028 0-0052
snvs_rtc 30370000.snvs:snvs-rtc-lp
```

Tip

This will list all RTCs including the non-I²C RTCs. Linux assigns RTC device IDs based on the device tree/aliases entries if present.

Date and time can be manipulated with the `hwclock` tool and the `date` command. To show the current date and time set on the target:

```
target:~$ date
Thu Jan  1 00:01:26 UTC 1970
```

Change the date and time with the `date` command. The `date` command sets the time with the following syntax “YYYY-MM-DD hh:mm:ss (+|-)hh:mm”:

```
target:~$ date -s "2022-03-02 11:15:00 +0100"
Wed Mar  2 10:15:00 UTC 2022
```

Note

Your timezone (in this example +0100) may vary.

Using the `date` command does not change the time and date of the RTC, so if we were to restart the target those changes would be discarded. To write to the RTC we need to use the `hwclock` command. Write the current date and time (set with the `date` command) to the RTC using the `hwclock` tool and reboot the target to check if the changes were applied to the RTC:

```
target:~$ hwclock -w
target:~$ reboot
.
.
.
target:~$ date
Wed Mar  2 10:34:06 UTC 2022
```

To set the time and date from the RTC use:

```
target:~$ date
Thu Jan  1 01:00:02 UTC 1970
target:~$ hwclock -s
target:~$ date
Wed Mar  2 10:45:01 UTC 2022
```

7.11.1 7.11.1 RTC Parameters

RTCs have a few abilities which can be read/set with the help of hwclock tool.

- We can check RTC supported features with:

```
target:~$ hwclock --param-get features
The RTC parameter 0x0 is set to 0x71.
```

What this value means is encoded in kernel, each set bit translates to:

```
#define RTC_FEATURE_ALARM           0
#define RTC_FEATURE_ALARM_RES_MINUTE 1
#define RTC_FEATURE_NEED_WEEK_DAY   2
#define RTC_FEATURE_ALARM_RES_2S    3
#define RTC_FEATURE_UPDATE_INTERRUPT 4
#define RTC_FEATURE_CORRECTION       5
#define RTC_FEATURE_BACKUP_SWITCH_MODE 6
#define RTC_FEATURE_ALARM_WAKEUP_ONLY 7
#define RTC_FEATURE_CNT              8
```

- We can check RTC BSM (Backup Switchover Mode) with:

```
target:~$ hwclock --param-get bsm
The RTC parameter 0x2 is set to 0x1.
```

- We can set RTC BSM with:

```
target:~$ hwclock --param-set bsm=0x2
The RTC parameter 0x2 will be set to 0x2.
```

What BSM values mean translates to these values:

```
#define RTC_BSM_DISABLED  0
#define RTC_BSM_DIRECT    1
#define RTC_BSM_LEVEL     2
#define RTC_BSM_STANDBY   3
```

Tip

You should set BSM mode to DSM or LSM for RTC to switch to backup power source when the initial power source is not available. Check **RV-3028** RTC datasheet to read what LSM (Level Switching Mode) and DSM (Direct Switching Mode) actually mean.

DT representation for I²C RTCs: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-fpsc-g-som.dtsi#L318>

And the additions on the expansion board: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip-peb-b-006.dtso#L150>

7.12 7.12 USB Host Controller

The USB controller of the i.MX 8M Plus SoC provides a low-cost connectivity solution for numerous consumer portable devices by providing a mechanism for data transfer between USB devices with a line/bus speed of up to 4 Gbit/s (SuperSpeed 'SS'). The USB subsystem has two independent USB controller cores.

The unified BSP includes support for mass storage devices and keyboards. Other USB-related device drivers must be enabled in the kernel configuration on demand. Due to udev, all mass storage devices connected get unique IDs and can be found in `/dev/disk/by-id`. These IDs can be used in `/etc/fstab` to mount the different USB memory devices in different ways.

In U-Boot USB mass storage devices will not automatically be detected on insertion. Instead USB devices will be probed once when USB is started.

```
u-boot=> usb start
```

If USB is already started and a new device is plugged in, USB needs to be restarted to trigger the device detection again.

```
u-boot=> usb reset
```

DT representation for USB Host: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip-peb-b-006.dtso#L193>

7.13 7.13 CAN FD

The phyVIP has two flexCAN interfaces supporting CAN FD. They are supported by the Linux standard CAN framework which builds upon then the Linux network layer. Using this framework, the CAN interfaces behave like an ordinary Linux network device, with some additional features special to CAN. More information can be found in the Linux Kernel documentation: <https://www.kernel.org/doc/html/latest/networking/can.html>

- Use:

```
target:~$ ip link
```

to see the state of the interfaces.

- To get information on fcan1, such as bit rate and error counters, type:

```
target:~$ ip -d -s link show fcan1
4: fcan1: <NOARP,UP,LOWER_UP,ECHO> mtu 72 qdisc pfifo_fast state UP mode DEFAULT group_
```

(continues on next page)

(continued from previous page)

```

↪default qlen 10
  link/can promiscuity 0 allmulti 0 minmtu 0 maxmtu 0
  can <FD> state ERROR-ACTIVE (berr-counter tx 0 rx 0) restart-ms 0
    bitrate 500000 sample-point 0.875
    tq 25 prop-seg 37 phase-seg1 32 phase-seg2 10 sjw 5 brp 1
    flexcan: tseg1 2..96 tseg2 2..32 sjw 1..16 brp 1..1024 brp_inc 1
    dbitrates 2000000 dsample-point 0.750
    dtq 25 dprop-seg 7 dphase-seg1 7 dphase-seg2 5 dsjw 2 dbrp 1
    flexcan: dtseg1 2..39 dtseg2 2..8 dsjw 1..4 dbrp 1..1024 dbrp_inc 1
    clock 40000000
    re-started bus-errors arbit-lost error-warn error-pass bus-off
      0      0      0      0      0      0      numtxqueues 1
↪numrxqueues 1 gso_max_size 65536 gso_max_segs 65535 tso_max_size 65536 tso_max_segs 65535
↪gro_max_size 65536
gso_ipv4_max_size 65536 gro_ipv4_max_size 65536 parentbus platform parentdev 443a0000.can
  RX:  bytes packets errors dropped missed  mcast
      0      0      0      0      0      0
  TX:  bytes packets errors dropped carrier collsns
      0      0      0      0      0      0

```

The output contains a standard set of parameters also shown for Ethernet interfaces, so not all of these are necessarily relevant for CAN (for example the MAC address). The following output parameters contain useful information:

fcan1	Interface Name
NOARP	CAN cannot use ARP protocol
MTU	Maximum Transfer Unit
RX packets	Number of Received Packets
TX packets	Number of Transmitted Packets
RX bytes	Number of Received Bytes
TX bytes	Number of Transmitted Bytes
errors...	Bus Error Statistics

The CAN configuration is done in the systemd configuration file `/lib/systemd/network/11-can.network`. For a persistent change of (as an example, the default bitrates), change the configuration in the BSP under `./meta-ampliphy/recipes-core/systemd/systemd-conf/11-can.network` in the root filesystem and rebuild the root filesystem.

Note

By default, we enable CAN-FD (flexible datarate) in our BSP. In case CAN Classic is required one needs to remove options `FDMODE` and `DataBitRate` from the `/lib/systemd/network/11-can.network` file.

To disable flexible datarate manually, one can use:

```

target:~$ ip link set fcan1 down
target:~$ ip link set fcan1 type can bitrate 500000 dbitrates 0 fd off
target:~$ ip link set fcan1 up

```

The bitrate can also be changed manually, for example:

```
target:~$ ip link set fcan1 down
target:~$ ip link set fcan1 txqueuelen 10 up type can bitrate 500000 sample-point 0.75 dbitrate
↪4000000 dsample-point 0.8 fd on
target:~$ ip link set fcan1 up
```

You can send messages with cansend or receive messages with candump:

```
target:~$ cansend fcan1 123#45.67
target:~$ candump fcan1
```

To generate random CAN traffic for testing purposes, use cangen:

```
target:~$ cangen
```

cansend --help and candump --help provide help messages for further information on options and usage.

Device Tree CAN configuration of imx8mp-phyflex-fpsc-g-som.dtsi: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-fpsc-g-som.dtsi#L121>

and of imx8mp-phyflex-phyvip-peb-b-006.dtso: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip-peb-b-006.dtso#L85>

7.14 7.14 Audio

Playback devices supported for phyVIP-i.MX 8M Plus are HDMI and the TI TLV320AIC3007 audio codec on the default PEB-B-006 expansion board. The connector X16 on the PEB-B-006 provides mono-speaker output, while the X17 connector provides input for the microphone. The PEB-B-006 also features a Class-D amplifier output on the X18 connector.

Note

For more details refer to the phyVIP-i.MX 8M Plus Hardware Manual and the Schematics.

To check if your soundcard driver is loaded correctly and what the device is called, type for playback devices:

```
target:~$ aplay -L
```

Or type for recording devices:

```
target:~$ arecord -L
```

7.14.1 7.14.1 Alsamixer

To inspect the capabilities of your soundcard, call:

```
target:~$ alsamixer
```

You should see a lot of options as the audio-IC has many features you can experiment with. It might be better to open alsamixer via ssh instead of the serial console, as the console graphical effects are better. You have either mono or stereo gain controls for all mix points. “MM” means the feature is muted (both output, left & right), which can be toggled by hitting ‘m’. You can also toggle by hitting ‘<’ for left and ‘>’ for right output. With the **tab** key, you can switch between controls for playback and recording.

7.14.2 7.14.2 Restore default volumes

There are some default settings stored in `/var/lib/alsa/asound.state`. You can save your current alsa settings with:

```
target:~$ alsactl --file </path/to/filename> store
```

You can restore saved alsa settings with:

```
target:~$ alsactl --file </path/to/filename> restore
```

7.14.3 7.14.3 ALSA configuration

Our BSP comes with a ALSA configuration file `/etc/asound.conf`.

The ALSA configuration file can be edited as desired or deleted since it is not required for ALSA to work properly.

```
target:~$ vi /etc/asound.conf
```

To set PEB-AV-10 as output, set *playback.pcm* from “dummy” to “pebav10”:

```
[...]

pcm.asymed {
    type asym
    playback.pcm "pebav10"
    capture.pcm "dsnoop"
}

[...]
```

If the sound is not audible change playback devices to the software volume control playback devices, set *playback.pcm* to the respective softvol playback device e.g. “softvol_pebav10”. Use alsamixer controls to vary the volume levels.

```
[...]

pcm.asymed {
    type asym
    playback.pcm "softvol_pebav10"
    capture.pcm "dsnoop"
}

[...]
```

7.14.4 7.14.4 Pulseaudio Configuration

For applications using *Pulseaudio*, check for available sinks:

```
target:~$ pactl list short sinks
```

To select the output device, type:

```
target:~$ pactl set-default-sink <sink_number>
```

7.14.5 7.14.5 Playback

Run speaker-test to check playback availability:

```
target:~$ speaker-test -c 2 -t wav
```

To playback simple audio streams, you can use aplay. For example to play the ALSA test sounds:

```
target:~$ aplay /usr/share/sounds/alsa/*
```

To playback other formats like mp3 for example, you can use Gstreamer:

```
target:~$ gst-launch-1.0 playbin uri=file:/path/to/file.mp3
```

7.14.6 7.14.6 Capture

arecord is a command-line tool for capturing audio streams which use Line In as the default input source. To select a different audio source you can use alsamixer. For example, switch on *Right PGA Mixer Mic3R* and *Left PGA Mixer Mic3R* in order to capture the audio from the microphone input of the TLV320-Codec using the 3.5mm jack.

```
target:~$ amixer -c "sndpebav10" sset 'Left PGA Mixer Mic3R' on
target:~$ amixer -c "sndpebav10" sset 'Right PGA Mixer Mic3R' on
```

```
target:~$ arecord -t wav -c 2 -r 44100 -f S16_LE test.wav
```

Hint

Since playback and capture share hardware interfaces, it is not possible to use different sampling rates and formats for simultaneous playback and capture operations.

Device Tree Audio configuration: <https://github.com/phytec/linux-phytec-imx/tree/v6.12.49-2.2.0-phy/arch/arm64/boot/dts/freescale/imx8mp-phyflex-phyvip-peb-b-006.dtso#L51>

7.15 7.15 Display

The phyVIP supports up to two different display outputs: HDMI and LVDS. Both can be used simultaneously, however availability of these interfaces depends on the extension board design.

Note

The default PEB-B-006 extension board only supports the HDMI interface.

The following table shows the required extensions and devicetree overlays for the different interfaces.

Interface	Expansion	devicetree overlay
HDMI	phyVIP	imx8mp-phyflex-phyvip-peb-b-006.dtbo (enabled by default)
LVDS0	not yet available	not yet available

7.15.1 7.15.1 Weston Configuration

In order to get an output from Weston on the correct display, it still needs to be configured correctly. This will be done at `/etc/xdg/weston/weston.ini`.

7.15.1.1 Single Display

In our BSP, the default Weston output is set to HDMI.

```
[output]
name=HDMI-A-1
mode=preferred

[output]
name=LVDS-1
mode=off
```

7.16 7.16 GPU

The i.MX 8M Plus has a Vivante GC7000UL GPU.

The phytec-vision-image has glmark2 as a GPU benchmark built in. To run the benchmark, execute:

```
target:~$ glmark2-es2-wayland
```

Note

For this benchmark to run properly, weston needs to be started successfully. For this display overlays might need to be loaded.

By default the benchmark is executed in 800x600 windowed mode. The benchmark can also be executed in fullscreen mode with the `--fullscreen` flag. Alternatively with the `--off-screen` flag the benchmark will not be shown on display.

7.17 7.17 Power Management

7.17.1 7.17.1 CPU Core Frequency Scaling

The CPU in the i.MX 8M Plus SoC is able to scale the clock frequency and the voltage. This is used to save power when the full performance of the CPU is not needed. Scaling the frequency and the voltage is referred to as ‘Dynamic Voltage and Frequency Scaling’ (DVFS). The i.MX 8M Plus BSP supports the DVFS feature. The Linux kernel provides a DVFS framework that allows each CPU core to have a min/max frequency and a governor that governs it. Depending on the i.MX 8 variant used, several different frequencies are supported.

Tip

Although the DVFS framework provides frequency settings for each CPU core, a change in the frequency settings of one CPU core always affects all other CPU cores too. So all CPU cores always share the same DVFS setting. An individual DVFS setting for each core is not possible.

- To get a complete list type:

```
target:~$ cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_available_frequencies
```

In case you have, for example, i.MX 8MPlus CPU with a maximum of approximately 1,6 GHz, the result will be:

```
1200000 1600000
```

- To ask for the current frequency type:

```
target:~$ cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_cur_freq
```

So-called governors are automatically selecting one of these frequencies in accordance with their goals.

- List all governors available with the following command:

```
target:~$ cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_available_governors
```

The result will be:

```
conservative ondemand userspace powersave performance schedutil
```

- **conservative** is much like the ondemand governor. It differs in behavior in that it gracefully increases and decreases the CPU speed rather than jumping to max speed the moment there is any load on the CPU.
- **ondemand** (default) switches between possible CPU core frequencies in reference to the current system load. When the system load increases above a specific limit, it increases the CPU core frequency immediately.
- **powersave** always selects the lowest possible CPU core frequency.
- **performance** always selects the highest possible CPU core frequency.
- **userspace** allows the user or userspace program running as root to set a specific frequency (e.g. to 1600000). Type:
- In order to ask for the current governor, type:

```
target:~$ cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_governor
```

You will normally get:

```
ondemand
```

- Switching over to another governor (e.g. userspace) is done with:

```
target:~$ echo userspace > /sys/devices/system/cpu/cpu0/cpufreq/scaling_governor
```

- Now you can set the speed:

```
target:~$ echo 1600000 > /sys/devices/system/cpu/cpu0/cpufreq/scaling_setspeed
```

For more detailed information about the governors, refer to the Linux kernel documentation in the linux kernel repository at [Documentation/admin-guide/pm/cpufreq.rst](#).

7.17.2 7.17.2 CPU Core Management

The i.MX 8M Plus SoC can have multiple processor cores on the die. The i.MX 8M Plus, for example, has 4 ARM Cores which can be turned on and off individually at runtime.

- To see all available cores in the system, execute:

```
target:~$ ls /sys/devices/system/cpu -l
```

- This will show, for example:

```
cpu0    cpu1    cpu2    cpu3    cpufreq
[...]
```

Here the system has four processor cores. By default, all available cores in the system are enabled to get maximum performance.

- To switch off a single-core, execute:

```
target:~$ echo 0 > /sys/devices/system/cpu/cpu3/online
```

As confirmation, you will see:

```
[ 110.505012] psci: CPU3 killed
```

Now the core is powered down and no more processes are scheduled on this core.

- You can use top to see a graphical overview of the cores and processes:

```
target:~$ htop
```

- To power up the core again, execute:

```
target:~$ echo 1 > /sys/devices/system/cpu/cpu3/online
```

7.18 7.18 Thermal Management

7.18.1 7.18.1 U-Boot

The previous temperature control in the U-Boot was not satisfactory. Now the u-boot has a temperature shutdown to prevent the board from getting too hot during booting. The temperatures at which the shutdown occurs are identical to those in the kernel.

The individual temperature ranges with the current temperature are displayed in the boot log:

```
CPU: Industrial temperature grade (-40C to 105C) at 33C
```

7.18.2 7.18.2 Kernel

The Linux kernel has integrated thermal management that is capable of monitoring SoC temperatures, reducing the CPU frequency, driving fans, advising other drivers to reduce the power consumption of devices, and – worst-case – shutting down the system gracefully (<https://www.kernel.org/doc/Documentation/thermal/sysfs-api.txt>).

This section describes how the thermal management kernel API is used for the i.MX 8M Plus SoC platform. The i.MX 8 has internal temperature sensors for the SoC.

- The current temperature can be read in millicelsius with:

```
target:~$ cat /sys/class/thermal/thermal_zone0/temp
```

- You will get, for example:

```
49000
```

There are two trip points registered by the imx_thermal kernel driver. These differ depending on the CPU variant. A distinction is made between Industrial and Commercial.

	Commercial	Industrial
passive (warning)	85°C	95°C
critical (shutdown)	90°C	100°C

(see kernel sysfs folder `/sys/class/thermal/thermal_zone0/`)

The kernel thermal management uses these trip points to trigger events and change the cooling behavior. The following thermal policies (also named thermal governors) are available in the kernel: Step Wise, Fair Share, Bang Bang, and Userspace. The default policy used in the BSP is `step_wise`. If the value of the SoC temperature in the sysfs file `temp` is above `trip_point_0`, the CPU frequency is set to the lowest CPU frequency. When the SoC temperature drops below `trip_point_0` again, the throttling is released.

Note

The actual values of the thermal trip points may differ since we mount CPUs with different temperature grades.

7.19 7.19 TPM

Some carrier boards/SOMs are equipped with a Trusted Platform Module (TPM) that provides hardware-based security functions.

Note

To check if TPM is supported on your hardware, refer to the hardware manual for your carrier board/SOM on the download page: <https://www.phytec.eu/en/produkte/embedded-vision/phyvip-vision-integration-platform-imx-8m-plus/#downloads>. Alternatively, check whether a TPM device file (`/dev/tpm*`) exists on your target system.

Here are some useful examples to work with the TPM

Generate 4-byte random value with TPM2 tools:

```
target:~$ tpm2_getrandom --hex 4
```

Generate 4-byte random value with OpenSSL tools:

```
target:~$ openssl rand -engine libtpm2tss --hex 4
```

Generate RSA private key and validate its contents:

```
target:~$ openssl genrsa -engine libtpm2tss -out /tmp/priv_key 512
Engine "tpm2tss" set.
target:~$ openssl rsa -check -in /tmp/priv_key -noout
RSA key ok
target:~$ cat /tmp/priv_key
-----BEGIN PRIVATE KEY-----
MIIBVQIBADANBgkqhkiG9w0BAQEFAASCAT8wggE7AgEAAkEAsvmcbxjwuKnYeuZ
2AVBmuLvYyqF/LpY0D3IB/v+YvEoLxdGGmjiFLECU6xZ1j3+dIt4Y1zbckS10cWT
I8mbSwIDAQABAKBoy8wrYNhmP/1kzUJIcLznPYJckGoZlFI1M7xjGSA9H1xDK6if
5g5CYCHPrbBp8e0mEokPRZoihxzGTxGPiahAiEA/70YM0pVZ5SD3YcRsWcQlkWI
M0SPUYg6vxxvGG9xp4FCIQDHB01RoHr+qXJwxIu3/3oQAUBI4ACJ4JRp0KelwhC0
LQIHANJzSvg/dak5l8pU55/99q3nbm7nPnnZSJiP0F6P62gjAiEAjf7qrFMF7Uyt
RkEjwbl2t5Z868FNARGGMVxZT4x+aF0CIGxlmP2pL8xFu1bWB282LSedqZUdQweL
Lxi7+svb2+uJ
-----END PRIVATE KEY-----
```

Note

Do NOT share your private RSA keys if you are going to use these keys for any security purposes.

Generate RSA public key and validate its contents:

```
target:~$ openssl rsa -in /tmp/test_key -pubout -out /tmp/pub_key
writing RSA key
target:~$ openssl pkey -inform PEM -pubin -in /tmp/pub_key -noout
target:~$ cat /tmp/pub_key
-----BEGIN PUBLIC KEY-----
MFwwDQYJKoZIhvcNAQEBBQADSwAwSAJBAMbL5nG8Y8Lip2HrmdgFQZri72Mqhfy6
WDg9yAf7/mLxKJcXRhpo4hSxAl0sWdY9/nSLeGnc23CktTnFkyPJm0sCAwEAAQ==
-----END PUBLIC KEY-----
```

7.20 Watchdog

The PHYTEC i.MX 8M Plus modules include a hardware watchdog that is able to reset the board when the system hangs. The watchdog is started on default in U-Boot with a timeout of 60s. So even during early kernel start, the watchdog is already up and running. The Linux kernel driver takes control over the watchdog and makes sure that it is fed. This section explains how to configure the watchdog in Linux using systemd to check for system hangs and during reboot.

7.20.1 7.20.1 Watchdog Support in systemd

Systemd has included hardware watchdog support since version 183.

- To activate watchdog support, the file `system.conf` in `/etc/systemd/` has to be adapted by enabling the options:

```
RuntimeWatchdogSec=60s
RebootWatchdogSec=10min
```

RuntimeWatchdogSec defines the timeout value of the watchdog, while *RebootWatchdogSec* defines the time-out when the system is rebooted. For more detailed information about hardware watchdogs under systemd can be found at <http://0pointer.de/blog/projects/watchdog.html>. The changes will take effect after a reboot or run:

```
target:~$ systemctl daemon-reload
```

7.21 7.21 Power Key

The X_ONOFF pin connected to the ON/OFF button can be pressed long (for 5 seconds) to trigger Power OFF without SW intervention or used to wake up the system out of suspend. With the `snvs_pwrkey` driver, the `KEY_POWER` event is also reported to userspace when the button is pressed. On default, systemd is configured to ignore such events. The function of Power OFF without SW intervention is not configured. Triggering a power off with systemd when pushing the ON/OFF button can be configured under `/usr/lib/systemd/logind.conf.d/00-systemd-conf-imx.conf` and set using:

```
HandlePowerKey=poweroff
```

To use the changes, reboot or restart the logind service.

```
systemctl restart systemd-logind
```

By default, ON/OFF button is configured to emit the `KEY_POWER` event code when short-pressed. This behavior can be verified using the `evtest` tool, which is included in our BSP:

```
target:~$ evtest /dev/input/by-path/platform-30370000.snvs:snvs-powerkey-event
Input driver version is 1.0.1
Input device ID: bus 0x19 vendor 0x0 product 0x0 version 0x0
Input device name: "30370000.snvs:snvs-powerkey"
Supported events:
Event type 0 (EV_SYN)
Event type 1 (EV_KEY)
    Event code 116 (KEY_POWER)
Properties:
Testing ... (interrupt to exit)
Event: time 1762862749.414004, type 1 (EV_KEY), code 116 (KEY_POWER), value 1
Event: time 1762862749.414004, ----- SYN_REPORT -----
Event: time 1762862749.542006, type 1 (EV_KEY), code 116 (KEY_POWER), value 0
Event: time 1762862749.542006, ----- SYN_REPORT -----
```

However, it is possible to reassign the short-press event by modifying the device-tree and changing the default `linux,keycode` property of the `snvs_pwrkey` node:

```
&snvs_pwrkey {
    linux,keycode = <KEY_SLEEP>;
};
```

In this example, the ON/OFF button is assigned KEY_SLEEP code, and pressing the button will now put the board into sleep mode by default.

All valid key codes are defined in the `include/uapi/linux/input-event-codes.h` file part of the Linux kernel.

7.22 7.22 AI Accelerators

In addition to the internal i.MX 8M Plus SoC NPU, phyVIP-i.MX 8M Plus supports an optional external AI accelerator connected via PCIe, depending on the hardware variant:

- **Hailo-8** NPU: up to 26 TOPS
- **Kinara ARA240** NPU: up to 40 TOPS

Note

Currently, only the Hailo-8 NPU is supported in the BSP. Kinara NPU support is not yet available. Please contact PHYTEC support for more information on driver availability.

7.22.1 7.22.1 Internal NPU

The internal i.MX 8M Plus NPU is accessed from Linux userspace via the VeriSilicon “galcore” kernel driver (shared with the Vivante GC7000UL GPU) and NXP’s eIQ TensorFlow Lite VX delegate (`libvx_delegate.so`).

Demo applications and a sample model are pre-installed under `/usr/bin/tensorflow-lite-2.19.0/examples`.

Sanity check – confirm the NPU device node is present:

```
target:~$ ls -l /dev/galcore
```

Benchmark and health report - compare CPU-only inference against NPU-delegated inference on the bundled sample model:

```
target:~$ cd /usr/bin/tensorflow-lite-2.19.0/examples
target:examples$ ./benchmark_model --graph=mobilenet_v1_1.0_224_quant.tflite
INFO: Created TensorFlow Lite XNNPACK delegate for CPU.
...
INFO: Inference timings in us: Init: 3991, First inference: 165064, Warmup
(avg): 160546, Inference (avg): 158634
```

```
target:examples$ ./benchmark_model --graph=mobilenet_v1_1.0_224_quant.tflite \
    --external_delegate_path=/usr/lib/libvx_delegate.so
INFO: External delegate path: [/usr/lib/libvx_delegate.so]
INFO: EXTERNAL delegate created.
...
INFO: Inference timings in us: Init: 15530, First inference: 3276819, Warmup
(avg): 3.27682e+06, Inference (avg): 3424.19
```

The *Inference (avg)* time drops from ~159 ms (CPU) to ~3.4 ms (NPU), roughly a 46x speedup, confirming the NPU is active and healthy.

For more information, see the [i.MX Machine Learning User's Guide](#).

7.22.2 Hailo-8 NPU

The Hailo-8 NPU is accessed from Linux userspace via the HailoRT (v4.24.0) software stack integrated in our BSP (hailortcli, hailo). The following examples show how to verify the accelerator and evaluate its performance.

Sanity check – confirm the device is detected:

```
target:~$ hailortcli scan
```

Identify the chip (firmware version, serial, part number):

```
target:~$ hailortcli fw-control identify
```

Benchmark throughput, latency and power on a compiled model (HEF):

```
target:~$ wget https://hailo-model-zoo.s3.eu-west-2.amazonaws.com/ModelZoo/Compiled/v2.18.0/
↪hailo8/resnet_v1_18.hef
target:~$ hailo run resnet_v1_18.hef
target:~$ hailortcli benchmark resnet_v1_18.hef
```

For more information, see the [HailoRT documentation](#).

7.23 ISP

The i.MX 8M Plus SoC contains an Image Signal Processor (ISP). For more information see Using the ISPs on the phyVIP i.MX 8M Plus documentation. This documentation is also available in German.

7.24 On-Chip OTP Controller (OCOTP_CTRL) - eFuses

The i.MX 8M Plus provides one-time programmable fuses to store information such as the MAC address, boot configuration, and other permanent settings (“On-Chip OTP Controller (OCOTP_CTRL)” in the i.MX 8M Plus Reference Manual). The following list is an abstract from the i.MX 8M Plus Reference Manual and includes some useful registers in the OCOTP_CTRL (at base address 0x30350000):

Name	Bank	Word	Memory offset at 0x30350000	Description
OCOTP_MAC_9	0	0	0x640	contains lower 32 bits of ENET0 MAC address
OCOTP_MAC_9	1	1	0x650	contains upper 16 bits of ENET0 MAC address and the lower 16 bits of ENET1 MAC address
OCOTP_MAC_9	2	2	0x660	contains upper 32 bits of ENET1 MAC address

A complete list and a detailed mapping between the fuses in the OCOTP_CTRL and the boot/mac/... configuration are available in the section “Fuse Map” of the i.MX 8M Plus Security Reference Manual.

7.24.1 7.24.1 Reading Fuse Values in uBoot

You can read the content of a fuse using memory-mapped shadow registers. To calculate the memory address, use the fuse Bank and Word in the following formula:

OCOTP_MAC_ADDR:

```
u-boot=> fuse read 9 0
```

7.24.2 7.24.2 Reading Fuse Values in Linux

To access the content of the fuses in Linux NXP provides the NVMEM_IMX_OCOTP module. All fuse content of the memory-mapped shadow registers is accessible via sysfs:

```
target:~$ hexdump /sys/devices/platform/soc@0/30000000.bus/30350000.efuse/imx-ocotp0/nvmem
```